arXiv:1511.02306v1 [quant-ph] 7 Nov 2015

# Quantum linear systems algorithm with exponentially improved dependence on precision

Andrew M. Childs* Robin Kothari† Rolando D. Somma‡

**Abstract**

Harrow, Hassidim, and Lloyd showed that for a suitably specified $N \times N$ matrix $A$ and $N$-dimensional vector $\vec{b}$, there is a quantum algorithm that outputs a quantum state proportional to the solution of the linear system of equations $A\vec{x} = \vec{b}$. If $A$ is sparse and well-conditioned, their algorithm runs in time $\mathrm{poly}(\log N, 1/\epsilon)$, where $\epsilon$ is the desired precision in the output state. We improve this to an algorithm whose running time is polynomial in $\log(1/\epsilon)$, exponentially improving the dependence on precision while keeping essentially the same dependence on other parameters. Our algorithm is based on a general technique for implementing any operator with a suitable Fourier or Chebyshev series representation. This allows us to bypass the quantum phase estimation algorithm, whose dependence on $\epsilon$ is prohibitive.

## 1 Introduction

Recently, Harrow, Hassidim, and Lloyd [HHL09] gave an efficient quantum algorithm for the Quantum Linear Systems Problem (QLSP). Here the goal is to prepare a quantum state $|x\rangle$ proportional to the solution $\vec{x}$ of a linear system of equations $A\vec{x} = \vec{b}$, given procedures for computing the entries of $A$ and for preparing a quantum state $|b\rangle$ proportional to $\vec{b}$. If the $N \times N$ matrix $A$ is sparse and well-conditioned, and if the procedures for computing entries of $A$ and for preparing $|b\rangle$ are efficient, then the Harrow–Hassidim–Lloyd (HHL) algorithm produces an $\epsilon$-approximation to the desired quantum state using $\mathrm{poly}(\log N, 1/\epsilon)$ resources (where the notation $\mathrm{poly}(m, n)$ denotes a function upper bounded by a polynomial in $m$ and $n$).

The core of the HHL algorithm is an efficient procedure for simulating the dynamics of quantum systems. Whereas straightforward approaches to quantum simulation using product formulas have complexity polynomial in $1/\epsilon$, recent work has given an algorithm with complexity $\mathrm{poly}(\log(1/\epsilon))$ [BCC+14], an exponential improvement in the dependence on $\epsilon$. However, the performance of the HHL algorithm is limited by its use of phase estimation, which requires $\Omega(1/\epsilon)$ uses of a unitary operation to estimate its eigenvalues to precision $\epsilon$. Thus, simply replacing the Hamiltonian simulation subroutine of the HHL algorithm with the best known method gives only a modest improvement, and in particular, still gives complexity $\mathrm{poly}(1/\epsilon)$.

In this work, we show how to circumvent the limitations of phase estimation, giving an algorithm for the QLSP that uses ideas from recent quantum simulation algorithms to apply the inverse of a matrix directly. Under the same assumptions as for the HHL algorithm, our algorithm uses $\mathrm{poly}(\log N, \log(1/\epsilon))$ resources, exponentially improving the dependence on the precision parameter.

*Department of Computer Science, Institute for Advanced Computer Studies, and Joint Center for Quantum Information and Computer Science, University of Maryland. amchilds@umd.edu

†Center for Theoretical Physics, Massachusetts Institute of Technology. rkothari@mit.edu

‡Theoretical Division, Los Alamos National Laboratory. somma@lanl.gov

To obtain this improvement, it is essential to consider QLSP as an inherently quantum problem, where the goal is to output a quantum state $|x\rangle$. Originally, the HHL algorithm was described as a method for sampling expectation values of $|x\rangle$, providing a classical output [HHL09]. For the expectation value version of the problem, sampling error alone rules out the possibility of an algorithm with complexity $\mathrm{poly}(\log(1/\epsilon))$, unless $\mathsf{BQP} = \mathsf{PP}$ [HHL09, Theorem 6]. However, the HHL algorithm actually solves the more general problem of outputting $|x\rangle$ (and the algorithm is commonly described in those terms [Har15]).

The improved performance of our approach may be especially useful when the quantum linear systems algorithm is used as a subroutine polynomially many times, so that its output must have inverse polynomial precision to guarantee that the final algorithm succeeds with high probability. An algorithm with $\mathrm{poly}(1/\epsilon)$ scaling incurs a polynomial overhead in running time due to error reduction, whereas an algorithm with $\mathrm{poly}(\log(1/\epsilon))$ scaling incurs only logarithmic overhead.

### 1.1 Problem statement

The QLSP can be stated more precisely as follows. We are given an $N \times N$ Hermitian matrix $A$ and a vector $\vec{b} = (b_1, b_2, \ldots, b_N)$. (The assumption that $A$ is Hermitian and square can be dropped without loss of generality [HHL09].) The problem is to create the quantum state $|x\rangle := \sum_{i=1}^{N} x_i|i\rangle / \|\sum_i x_i|i\rangle\|$, where the vector $\vec{x} = (x_1, x_2, \ldots, x_N)$ is defined by the equation $A\vec{x} = \vec{b}$. To obtain an algorithm running in time $\mathrm{poly}(\log N)$, we require succinct representations of $A$ and $\vec{b}$. As in [HHL09], we assume that access to $A$ and $\vec{b}$ is provided by black-box subroutines. For the vector $\vec{b}$, we assume there is a procedure $\mathcal{P}_B$ that produces the quantum state $|b\rangle := \sum_i b_i|i\rangle / \|\sum_i b_i|i\rangle\|$. For the matrix $A$, we assume there is a procedure $\mathcal{P}_A$ that computes the locations and values of the nonzero entries. Specifically, as in the best known algorithm for Hamiltonian simulation [BCK15], we assume $\mathcal{P}_A$ allows us to perform the map

$$|j, \ell\rangle \mapsto |j, \nu(j, \ell)\rangle \tag{1}$$

for any $j \in [N] := \{1, \ldots, N\}$ and $\ell \in [d]$, where $d$ is the maximum number of nonzero entries in any row or column, and $\nu \colon [N] \times [d] \to [N]$ computes the row index of the $\ell^{\text{th}}$ nonzero entry of the $j^{\text{th}}$ column. The procedure $\mathcal{P}_A$ also allows us to perform the map

$$|j, k, z\rangle \mapsto |j, k, z \oplus A_{jk}\rangle \tag{2}$$

for any $j, k \in [N]$, where the third register holds a bit string representing an entry of $A$. We assume the entries of $A$ can be represented exactly (or to sufficiently high precision that any error can be neglected).

Note that for the map (1), we assume that the locations of the nonzero entries of $A$ can be computed in place, as in previous work on Hamiltonian simulation [BC12, BCK15]. This is possible if we can efficiently compute both $(j, \ell) \mapsto \nu(j, \ell)$ and the reverse map $(j, \nu(j, \ell)) \mapsto \ell$, which is possible for typical implicitly specified matrices $A$. Alternatively, if $\nu$ provides the nonzero entries in ascending order, we can compute the reverse map with only a $\log d$ overhead by binary search. In the worst case, if the entries are unordered, there may be an additional factor of $O(\sqrt{d})$ using Grover's algorithm [Gro96] to compute the reverse map.

While the HHL algorithm solves the QLSP for all such matrices $A$, it is efficient only when $A$ is sparse and well-conditioned. (We discuss later how the sparsity assumption can be slightly relaxed.) An $N \times N$ matrix is called *$d$-sparse* if it has at most $d$ nonzero entries in any row or column. We call it simply *sparse* if $d = \mathrm{poly}(\log N)$. We call a matrix well-conditioned if its condition number is $\mathrm{poly}(\log N)$, where the *condition number* of a matrix is the ratio of the largest to smallest singular

value, and undefined when the smallest singular value of $A$ is 0 (i.e., when $A$ is not invertible). Since $A$ is Hermitian, its singular values and eigenvalues are equal in magnitude.

It will be convenient to quantify the resource requirements of algorithms solving the QLSP using two measures. First, by *query complexity* we mean the number of uses of the procedure $\mathcal{P}_A$ (treating this procedure as a black box). By *gate complexity*, we mean the total number of 2-qubit gates used in the algorithm. We say an algorithm is *gate-efficient* if its gate complexity is larger than its query complexity only by logarithmic factors. Formally, an algorithm with query complexity $Q$ is gate-efficient if its gate complexity is $O(Q \operatorname{poly}(\log Q, \log N))$.

Formally, we define the Quantum Linear Systems Problem as follows:

**Problem 1** (QLSP)**.** Let $A$ be an $N \times N$ Hermitian matrix with condition number $\kappa$, $\|A\| = 1$, and at most $d$ nonzero entries in any row or column. Let $\vec{b}$ be an $N$-dimensional vector, and let $\vec{x} := A^{-1}\vec{b}$. We define the quantum states $|b\rangle$ and $|x\rangle$ as

$$|b\rangle := \frac{\sum_i b_i|i\rangle}{\|\sum_i b_i|i\rangle\|} \quad \text{and} \quad |x\rangle := \frac{\sum_i x_i|i\rangle}{\|\sum_i x_i|i\rangle\|}. \tag{3}$$

Given access to a procedure $\mathcal{P}_A$ that computes entries of $A$ as described in equations (1) and (2) and a procedure $\mathcal{P}_B$ that prepares the state $|b\rangle$ in time $O(\operatorname{poly}(\log N))$, the goal is to output a state $|\tilde{x}\rangle$ such that $\||\tilde{x}\rangle - |x\rangle\| \leq \epsilon$, succeeding with probability $\Omega(1)$ (say, at least $1/2$), with a flag indicating success.

Note that while we only demand success with bounded error, repeating the procedure until it is successful gives an algorithm that is always correct and whose expected running time is asymptotically the same. Alternatively, by repeating the procedure $O(\log(1/\epsilon))$ times, we can give an algorithm that always outputs a state $\epsilon$-close to the desired one. To achieve this, the running time is simply multiplied by a factor of $O(\log(1/\epsilon))$.

### 1.2 Results

Harrow, Hassidim, and Lloyd [HHL09] present an algorithm for the QLSP that is efficient when $A$ is sparse and well-conditioned (i.e., when $d$ and $\kappa$ are both $\operatorname{poly}(\log N)$).

**Theorem 2** (HHL algorithm)**.** *The QLSP can be solved by a gate-efficient algorithm that makes* $O\big(\frac{d\kappa^2}{\epsilon} \operatorname{poly}(\log(d\kappa/\epsilon))\big)$ *queries to the oracle* $\mathcal{P}_A$ *and* $O(d\kappa \operatorname{poly}(\log(d\kappa/\epsilon)))$ *uses of* $\mathcal{P}_B$*.*

The stated complexity uses the best known results on Hamiltonian simulation [BCK15], improving the $d$-dependence compared to [HHL09]. More generally, the HHL algorithm also works assuming only the ability to efficiently solve the Hamiltonian simulation problem for $A$, i.e., to efficiently implement the unitary operation $\exp(-iAt)$, without having direct access to $\mathcal{P}_A$. In other words, the HHL algorithm uses Hamiltonian simulation for $A$ as a black box.

We improve the $\epsilon$-dependence of the HHL algorithm from $\operatorname{poly}(1/\epsilon)$ to $\operatorname{poly}(\log(1/\epsilon))$, keeping essentially the same dependence on the other parameters. We provide two algorithms for the QLSP, one based on decomposing an operator using its Fourier series and another based on a decomposition into Chebyshev polynomials. The Fourier approach has slightly worse dependence on $\log(1/\epsilon)$, but uses Hamiltonian simulation as a black box only and is therefore more generally applicable.

**Theorem 3** (Fourier approach)**.** *The QLSP can be solved with* $O(\kappa\sqrt{\log(\kappa/\epsilon)})$ *uses of a Hamiltonian simulation algorithm that approximates* $\exp(-iAt)$ *for* $t = O(\kappa \log(\kappa/\epsilon))$ *with precision* $O(\epsilon/(\kappa\sqrt{\log(\kappa/\epsilon)}))$*. Using the best known algorithm for Hamiltonian simulation [BCK15], this makes* $O(d\kappa^2 \log^{2.5}(\kappa/\epsilon))$ *queries to* $\mathcal{P}_A$*, makes* $O(\kappa\sqrt{\log(\kappa/\epsilon)})$ *uses of* $\mathcal{P}_B$*, and has gate complexity* $O(d\kappa^2 \log^{2.5}(\kappa/\epsilon)(\log N + \log^{2.5}(\kappa/\epsilon)))$*.*

The second approach uses the oracle for the entries of $A$ directly without using Hamiltonian simulation as a subroutine, achieving better dependence on $\epsilon$.

**Theorem 4** (Chebyshev approach)**.** *The QLSP can be solved using $O(d\kappa^2 \log^2(d\kappa/\epsilon))$ queries to $\mathcal{P}_A$ and $O(\kappa \log(d\kappa/\epsilon))$ uses of $\mathcal{P}_B$, with gate complexity $O(d\kappa^2 \log^2(d\kappa/\epsilon)(\log N + \log^{2.5}(d\kappa/\epsilon)))$.*

Both our algorithms achieve $\mathrm{poly}(\log(1/\epsilon))$ dependence on error, but are incomparable. The Fourier approach is more general, applying whenever the Hamiltonian $A$ can be efficiently simulated (even if it is not necessarily sparse). The Chebyshev approach is more efficient but applies only to sparse Hamiltonians.

Ambainis later improved the $\kappa$-dependence of the HHL algorithm from quadratic to nearly linear [Amb12], which is essentially optimal since the dependence on $\kappa$ cannot be made sublinear [HHL09]. Since Ambainis's approach crucially uses phase estimation, applying it to our algorithms would increase their $\epsilon$-dependence back to $\mathrm{poly}(1/\epsilon)$. By carefully modifying the technique, we simultaneously achieve nearly linear scaling in $\kappa$ and logarithmic dependence on $\epsilon$.

**Theorem 5** (Linear scaling in $\kappa$)**.** *The QLSP can be solved by a gate-efficient algorithm that makes $O\big(d\kappa\,\mathrm{poly}(\log(d\kappa/\epsilon))\big)$ queries to the oracles $\mathcal{P}_A$ and $\mathcal{P}_B$.*

Additionally, if $\vec{b}$ is known to lie in an invariant subspace of $A$ of condition number $\kappa' < \kappa$, then $\kappa$ can be replaced by $\kappa'$ in any of our upper bounds (as in the HHL algorithm). In particular, this means our algorithm works even when $A$ is not invertible as long as $\vec{b}$ is known to lie outside the null space of $A$, i.e., when $\vec{x} = A^{-1}\vec{b}$ is well defined. This property is useful when the matrix $A$ is rectangular, as the reduction in [HHL09] from rectangular to square matrices produces noninvertible matrices.

### 1.3 High-level overview

We now provide a high-level overview of our approach. The QLSP is equivalent to applying the (non-unitary) operator $A^{-1}$ to the state $|b\rangle$. Our general strategy is to represent $A^{-1}$, the operator we would like to perform, as a linear combination of unitaries we know how to perform. We then show how such a representation allows us to implement $A^{-1}$.

Our technique for implementing linear combinations of unitary operations arises from previous work on Hamiltonian simulation [BCC+14, Kot14, BCC+15, BCK15] (see also [SOG+02, CW12] for previous related approaches). As an example, consider implementing the operator $M = U_0 + U_1$, where $U_0$ and $U_1$ are unitaries with known quantum circuits. To implement $M$ on a state $|\psi\rangle$, we start with the state $|+\rangle|\psi\rangle$, where $|\pm\rangle := \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$. We then perform the unitary $|0\rangle\langle 0| \otimes U_0 + |1\rangle\langle 1| \otimes U_1$, giving the state $\frac{1}{\sqrt{2}}(|0\rangle U_0|\psi\rangle + |1\rangle U_1|\psi\rangle)$. If we measure the first qubit in the $\{|+\rangle, |-\rangle\}$ basis and obtain the $|+\rangle$ outcome, then we prepare a state proportional to $M|\psi\rangle$. If we have the ability to create multiple copies of $|\psi\rangle$ or reflect about $|\psi\rangle$, then we can create the output state with high probability by repeating this process until we get the desired measurement outcome or by using amplitude amplification. As we describe in Section 2, this strategy for implementing a linear combination of unitaries works more generally.

However, it is unclear how to decompose $A^{-1}$ as a linear combination of easy-to-implement unitaries. Such a decomposition depends on what unitaries are used as elementary building blocks. Our first choice is to use the unitaries $\exp(-iAt)$, which can be implemented using any Hamiltonian simulation algorithm. We then have to represent $A^{-1}$ as $\sum_j \alpha_j \exp(-iAt_j)$ for some coefficients $\alpha_j$ and evolution times $t_j$. Since both sides of the equation are diagonal in the same basis, this is equivalent to representing $x^{-1}$ as a linear combination of $\sum_j \alpha_j \exp(-ixt_j)$. The representation

only needs to be correct for $x \in D_\kappa := [-1, -1/\kappa] \cup [1/\kappa, 1]$ since we know the eigenvalues of $A$ fall in that range. For this range of $x$, we show how to approximate $x^{-1}$ as a linear combination of these unitaries in Section 3. Our strategy for doing this broadly comprises the following steps. Since $x^{-1}$ is unbounded at the origin, we first "tame" the function by multiplying it with a function that is close to 1 in the domain we care about but 0 near the origin, so that the overall function is bounded. We then perform the Fourier transform to obtain an integral over various $\exp(-ixt)$. After making some approximations, we discretize the integral to obtain a finite sum over $\exp(-ixt)$.

Instead of using the unitaries $e^{-iAt}$, our second approach uses the operators $\mathcal{T}_n(A/d)$ as its building blocks, where $d$ is the sparsity of $A$ and $\mathcal{T}_n$ is the $n^{\text{th}}$ Chebyshev polynomial of the first kind. These operators can be efficiently implemented using quantum walks [Chi10, BC12]. Now the problem is to represent $x^{-1}$ as a linear combination $\sum_n \alpha_n T_n(x/d)$. Our strategy for obtaining this representation is similar to the previous case. We first tame the function by making it bounded near the origin. We thereby obtain a function that is exactly representable as a linear combination of Chebyshev polynomials, and we approximate this function by dropping low-weight terms, as described in Section 4. Combining these decompositions with the linear combination of unitaries strategy outlined above yields our main results, Theorem 3 and Theorem 4.

Lastly, we show how to decrease the $\kappa$-dependence of our algorithms to nearly linear while retaining the desired $\text{poly}(\log(1/\epsilon))$ dependence on $\epsilon$. This improvement uses the observation that the complexity of our algorithms (and the HHL algorithm) is a product of two terms that depend on the eigenvalue $\lambda$ of $A$. Whereas the product of the maximum values of these two terms is quadratic in $\kappa$, the maximum of the product is only linear in $\kappa$. By introducing a technique called variable-time amplitude amplification, Ambainis exploited this observation to improve the $\kappa$-dependence of the HHL algorithm [Amb12]. Since our approach deliberately avoids phase estimation, it is not immediately clear how to invoke a strategy based on resolving the eigenvalues of $A$. Nevertheless, we show that a careful application of low-precision phase estimation suffices to resolve the eigenvalues into buckets of exponentially increasing size, which allows us to apply the appropriate approximation of $A^{-1}$ for a given eigenvalue range. We then apply variable-time amplitude amplification on this algorithm, followed by additional post-processing to remove information left behind by phase estimation. In Section 5 we establish Theorem 5 and show how the $\kappa$-dependence of both algorithms can be made nearly linear, just as Ambainis showed for the original HHL algorithm.

## 2 Implementing a linear combination of unitaries

We start by showing in Section 2.1 how to implement an operation $M$ that can be expressed as a linear combination of implementable unitaries. We then explain in Section 2.2 how this primitive can be applied to solve the QLSP, given a suitable decomposition of $A^{-1}$.

### 2.1 Framework

A technique for implementing linear combinations of unitaries was introduced in some recent Hamiltonian simulation algorithms [BCC+15, BCK15]. Since quantum simulation is unitary, $M$ is (nearly) unitary in the simulation algorithms based on this technique. Furthermore, in Hamiltonian simulation we only have one copy of the input state. Under these circumstances, one can use oblivious amplitude amplification [BCC+14] to implement $M$. In the QLSP, we can create multiple copies of the input state $|b\rangle$, so we do not require a tool like oblivious amplitude amplification (which would not work anyway when $M$ is far from unitary).

More precisely, our technique for implementing a linear combination of unitary operations is as follows. Let $M = \sum_i \alpha_i U_i$ be a linear combination of unitary matrices $U_i$ with $\alpha_i > 0$. We

assume $\alpha_i > 0$ without loss of generality since a phase can be subsumed into $U_i$. We show how to implement $M$ probabilistically using unitary operations $U$ and $V$ defined as follows. The operation $U := \sum_i |i\rangle\langle i| \otimes U_i$ implements $U_i$ conditioned on the value of a control register. The operation $V$ maps $|0^m\rangle$ to $\frac{1}{\sqrt{\alpha}} \sum_i \sqrt{\alpha_i}|i\rangle$, where $\alpha := \|\vec{\alpha}\|_1 = \sum_i \alpha_i$. Then, as shown in [Kot14, Lemma 2.1], we can implement $M$ in the following sense.

**Lemma 6.** *Let $M = \sum_i \alpha_i U_i$ be a linear combination of unitaries $U_i$ with $\alpha_i > 0$ for all $i$. Let $V$ be any operator that satisfies $V|0^m\rangle = \frac{1}{\sqrt{\alpha}} \sum_i \sqrt{\alpha_i}|i\rangle$, where $\alpha := \sum_i \alpha_i$. Then $W := V^\dagger U V$ satisfies*

$$W|0^m\rangle|\psi\rangle = \frac{1}{\alpha}|0^m\rangle M|\psi\rangle + |\Psi^\perp\rangle \tag{4}$$

*for all states $|\psi\rangle$, where $U := \sum_i |i\rangle\langle i| \otimes U_i$ and $(|0^m\rangle\langle 0^m| \otimes \mathbb{1})|\Psi^\perp\rangle = 0$.*

The lemma can be generalized to the case where $M = \sum_i \alpha_i T_i$, where each operator $T_i$ is not necessarily unitary, but is a block of a unitary, i.e., there exists a $U_i$ for which $U_i|0^t\rangle|\phi\rangle = |0^t\rangle T_i|\phi\rangle + |\Phi^\perp\rangle$ for all states $|\phi\rangle$, where $t \geq 0$ is an integer and $|\Phi^\perp\rangle$ has no overlap on states with $|0^t\rangle$ in the first register. We consider this more general situation in Lemma 7 below.

In Lemma 6, the operator $W$ can be thought of as a postselected or probabilistic implementation of $M$ in the sense that, if we measure the first $m$ qubits of $W|0^m\rangle|\psi\rangle$ and observe the output $|0^m\rangle$, the state of the second register is proportional to $M|\psi\rangle$. This successful outcome occurs with probability $(\|M|\psi\rangle\|/\alpha)^2$.

In our application, since we have the ability to create copies of the input state, we can repeat this process $O((\alpha/\|M|\psi\rangle\|)^2)$ times until the measurement yields the desired outcome. Alternately, we can construct the state with high probability using amplitude amplification [BHMT02]. Amplitude amplification requires us to reflect about the starting state $|\psi\rangle$, which in our application is $|b\rangle$. With two uses of the procedure $\mathcal{P}_B$ for preparing $|b\rangle$ (one performed in reverse), we can reflect about $|b\rangle$. Since amplitude amplification yields a quadratic speedup, we obtain the desired state after $O(\alpha/\|M|\psi\rangle\|)$ rounds. Combining amplitude amplification and a more general Lemma 6, we get the following procedure for implementing (submatrices of) linear combinations of unitary operations.

**Lemma 7** (Non-Unitary LCU Lemma)**.** *Let $M = \sum_i \alpha_i T_i$ with $\alpha_i > 0$ for some (not necessarily unitary) operators $\{T_i\}$. Let $\{U_i\}$ be a set of unitaries such that*

$$U_i|0^t\rangle|\phi\rangle = |0^t\rangle T_i|\phi\rangle + |\Phi_i^\perp\rangle \tag{5}$$

*for all states $|\phi\rangle$, where $t$ is a nonnegative integer and $(|0^t\rangle\langle 0^t| \otimes \mathbb{1})|\Phi_i^\perp\rangle = 0$. Given an algorithm $\mathcal{P}_B$ for creating a state $|b\rangle$, there is a quantum algorithm that exactly prepares the quantum state $M|b\rangle/\|M|b\rangle\|$ with constant success probability making $O(\alpha/\|M|b\rangle\|)$ uses of $\mathcal{P}_B$, $U$, and $V$, where*

$$U := \sum_i |i\rangle\langle i| \otimes U_i, \quad V|0^m\rangle = \frac{1}{\sqrt{\alpha}} \sum_i \sqrt{\alpha_i}|i\rangle, \quad \text{and} \quad \alpha := \sum_i \alpha_i, \tag{6}$$

*and that outputs a bit indicating whether it was successful.*

*Proof.* We start by implementing the linear combination $M' = \sum_i \alpha_i U_i$. Using Lemma 6, we have

$$W|0^m\rangle|\psi\rangle = \frac{1}{\alpha}|0^m\rangle M'|\psi\rangle + |\Psi^\perp\rangle, \tag{7}$$

where $W = V^\dagger U V$. Now consider the action of $W$ on $|\psi\rangle = |0^t\rangle|\phi\rangle$:

$$W|0^{m+t}\rangle|\phi\rangle = \frac{1}{\alpha}|0^m\rangle\Big(\sum_i \alpha_i U_i\Big)|0^t\rangle|\phi\rangle + |\Psi^\perp\rangle \tag{8}$$

$$= \frac{1}{\alpha}|0^m\rangle|0^t\rangle\Big(\sum_i \alpha_i T_i\Big)|\phi\rangle + \frac{1}{\alpha}|0^m\rangle\Big(\sum_i \alpha_i |\Phi_i^\perp\rangle\Big) + |\Psi^\perp\rangle \tag{9}$$

$$= \frac{1}{\alpha}|0^{m+t}\rangle M|\phi\rangle + |\Xi^\perp\rangle, \tag{10}$$

where $|\Xi^\perp\rangle$ satisfies $(|0^{m+t}\rangle\langle 0^{m+t}| \otimes \mathbb{1})|\Xi^\perp\rangle = 0$.

Now since we want to create a state proportional to $M|b\rangle$, let us plug in $|\phi\rangle = |b\rangle$. For convenience, let $r = m + t$, which gives

$$W|0^r\rangle|b\rangle = \frac{1}{\alpha}|0^r\rangle M|b\rangle + |\Xi^\perp\rangle = \Big(\frac{1}{\alpha}\|M|b\rangle\|\Big)|0^r\rangle\frac{M|b\rangle}{\|M|b\rangle\|} + |\Xi^\perp\rangle. \tag{11}$$

In words, applying $W$ on $|0^r\rangle|b\rangle$ followed by a measurement on the first $r$ qubits will yield the desired state $M|b\rangle/\|M|b\rangle\|$ with probability $(\|M|b\rangle\|/\alpha)^2$.

Since $\mathcal{P}_B$ is an algorithm that creates the state $|b\rangle$, we can also use it to reflect about the state $|0^r\rangle|b\rangle$. Specifically, say $\mathcal{P}_B$ performs the map $\mathcal{P}_B|0^s\rangle = |b\rangle$, then the operator $\mathcal{P}_B(\mathbb{1} - 2|0^{r+s}\rangle\langle 0^{r+s}|)\mathcal{P}_B^\dagger$ reflects about the state $|0^r\rangle|b\rangle$.

Given an algorithm that creates a desired state from an initial state with probability $q$, and the ability to reflect about the initial state, amplitude amplification [BHMT02] is a procedure for creating the desired state with probability (say) $1/2$ that uses $O(1/\sqrt{q})$ uses of the algorithm and the reflection map. Note that amplitude amplification also works when the probability $q$ is unknown, and hence we do not need an estimate of $\|M|b\rangle\|$ before we begin.

Using amplitude amplification we get an algorithm that creates the quantum state $M|b\rangle/\|M|b\rangle\|$ and makes $O(\alpha/\|M|b\rangle\|)$ uses of $W = V^\dagger U V$ and $\mathcal{P}_B$. □

To be completely general, we have stated our results for an arbitrary set of unitaries $\{U_i\}$ and quantified costs in terms of uses of $U := \sum_i |i\rangle\langle i| \otimes U_i$. When we apply these results, we will need to compute the complexity of implementing $U$ in terms of the complexities of implementing the $U_i$. The query complexity of implementing $U$ is precisely the maximum query complexity of implementing any of the $U_i$, which is optimal since $U$ cannot be easier to implement than any particular $U_i$.

On the other hand, the gate complexity of $U$ is not easy to characterize in terms of the gate complexities of $U_i$. In general, the different $U_i$ may be completely unrelated and the cost of implementing $U$ may be greater than the gate cost of the most expensive $U_i$. However, in our applications the matrices $\{U_i\}$ are related and are, in fact, powers of a single unitary $Y$. In this case the gate complexity of implementing $U$ behaves nicely, as we show below.

**Lemma 8.** *Let $U = \sum_{i=0}^N |i\rangle\langle i| \otimes Y^i$, where $Y$ is a unitary with gate complexity $G$. Then the gate complexity of $U$ is $O(NG)$.*

*Proof.* Let $n := \lfloor \log N \rfloor$ and consider the unitary $Y^{2^j}$ for $j \in \{0, \ldots, n\}$, which has gate complexity at most $2^j G$. Hence the controlled version of this unitary, c-$Y^{2^j}$, controlled by a single qubit, has gate complexity $O(2^j G)$. The unitary $U$ can then be implemented by a circuit that performs, for all $j \in \{0, \ldots, n\}$, the operation c-$Y^{2^j}$ on the second register controlled by the $j^{\text{th}}$ qubit of the first register $|i\rangle$. If the first register is in state $|i\rangle$, the operation performed on the second register is exactly $Y^i$, due to the binary encoding of the integer $i$. The gate complexity of this circuit is $O(\sum_j 2^j G) = O(NG)$. □

We can implement gates of the form $U = \sum_{i=0}^{N}\sum_{j=0}^{M} |i,j\rangle\langle i,j| \otimes Y^{ij}$ similarly. We use an additional register to compute the product of $i$ and $j$ and, conditional on this register, we apply $Y^{ij}$ as described in Lemma 8. Last, we uncompute the product in the additional register. The gate complexity in this case is $O(NMG)$.

### 2.2 Application to QLSP

We now describe how this technique can be applied to solve the QLSP. Suppose we can approximate $A^{-1}$ by a linear combination of operators $T_i$ that are either unitary or can be implemented by unitaries as in (5). Then we can implement this linear combination of operators, which is approximately $A^{-1}$, using Lemma 7. The following sections, Section 3 and Section 4, establish that $A^{-1}$ can indeed be represented in such a way using two different choices for the operators $T_i$ that correspond to using a Fourier decomposition and a Chebyshev decomposition, respectively.

For our application, we need to show that implementing an operator close to $A^{-1}$ yields a state close to the desired one. More precisely, we need to show that if two operators $C$ and $D$ are close, then the normalized states $C|\psi\rangle/\|C|\psi\rangle\|$ and $D|\psi\rangle/\|D|\psi\rangle\|$ are also close.

**Proposition 9.** *Let $C$ be a Hermitian operator with $\|C^{-1}\| \leq 1$ (i.e., the smallest eigenvalue of $C$ in absolute value is at least 1) and let $D$ be an operator that satisfies $\|C - D\| \leq \epsilon < 1/2$. Then the states $|x\rangle := C|\psi\rangle/\|C|\psi\rangle\|$ and $|\tilde{x}\rangle := D|\psi\rangle/\|D|\psi\rangle\|$ satisfy $\||x\rangle - |\tilde{x}\rangle\| < 4\epsilon$.*

*Proof.* Without loss of generality we assume $\||\psi\rangle\| = 1$. Using the triangle inequality, we get

$$\||x\rangle - |\tilde{x}\rangle\| = \left\| \frac{C|\psi\rangle}{\|C|\psi\rangle\|} - \frac{D|\psi\rangle}{\|D|\psi\rangle\|} \right\| \leq \left\| \frac{C|\psi\rangle}{\|C|\psi\rangle\|} - \frac{C|\psi\rangle}{\|D|\psi\rangle\|} \right\| + \left\| \frac{C|\psi\rangle}{\|D|\psi\rangle\|} - \frac{D|\psi\rangle}{\|D|\psi\rangle\|} \right\|. \tag{12}$$

Again using the triangle inequality, we have $\|C|\psi\rangle\| \leq \|D|\psi\rangle\| + \|(C-D)|\psi\rangle\| \leq \|D|\psi\rangle\| + \epsilon$, which yields

$$\big|\|D|\psi\rangle\| - \|C|\psi\rangle\|\big| \leq \epsilon \qquad \text{and} \qquad \|D|\psi\rangle\| \geq \|C|\psi\rangle\| - \epsilon \geq 1 - \epsilon, \tag{13}$$

where we used the fact that $\|C|\psi\rangle\| \geq 1$ since the smallest eigenvalue of $C$ in absolute value is at least 1. Then we can upper bound the first term of the right-hand side of (12) as follows:

$$\left\| \frac{C|\psi\rangle}{\|C|\psi\rangle\|} - \frac{C|\psi\rangle}{\|D|\psi\rangle\|} \right\| \leq \frac{\big|\|D|\psi\rangle\| - \|C|\psi\rangle\|\big|}{\|D|\psi\rangle\|} \leq \frac{\epsilon}{\|D|\psi\rangle\|} \leq \frac{\epsilon}{1-\epsilon} < 2\epsilon. \tag{14}$$

Analogously, we can bound the second term on the right-hand side of (12) as follows:

$$\left\| \frac{C|\psi\rangle}{\|D|\psi\rangle\|} - \frac{D|\psi\rangle}{\|D|\psi\rangle\|} \right\| \leq \frac{\big|\|C|\psi\rangle\| - \|D|\psi\rangle\|\big|}{\|D|\psi\rangle\|} \leq \frac{\epsilon}{\|D|\psi\rangle\|} \leq \frac{\epsilon}{1-\epsilon} < 2\epsilon. \tag{15}$$

Since both terms are at most $2\epsilon$, we have $\||x\rangle - |\tilde{x}\rangle\| < 4\epsilon$. □

Finally, we state a corollary that captures how we apply Lemma 7 to implement $A^{-1}$. This allows us to focus on approximating $1/x$ by a suitable linear combination of other functions. We say that functions $f$ and $g$ are $\epsilon$-close on a domain $D \subseteq \mathbb{R}$ if $|f(x) - g(x)| \leq \epsilon$ for all $x \in D$.

**Corollary 10.** *Let $A$ be a Hermitian operator with eigenvalues in a domain $D \subseteq \mathbb{R}$. Suppose the function $f : D \to \mathbb{R}$ satisfies $|f(x)| \geq 1$ for all $x \in D$ and is $\epsilon$-close to $\sum_i \alpha_i T_i$ on $D$ for some coefficients $\alpha_i > 0$ and functions $T_i \colon D \to \mathbb{R}$. Let $\{U_i\}$ be a set of unitaries such that*

$$U_i|0^t\rangle|\phi\rangle = |0^t\rangle T_i(A)|\phi\rangle + |\Phi_i^{\perp}\rangle \tag{16}$$

*for all states $|\phi\rangle$, where $t$ is a nonnegative integer and $(|0^t\rangle\langle 0^t| \otimes \mathbb{1})|\Phi_i^\perp\rangle = 0$. Given an algorithm $\mathcal{P}_B$ for creating a state $|b\rangle$, there is a quantum algorithm that prepares a quantum state $4\epsilon$-close to $f(A)|b\rangle/\|f(A)|b\rangle\|$, succeeding with constant probability, that makes $O(\alpha/\|f(A)|b\rangle\|) = O(\alpha)$ uses of $\mathcal{P}_B$, $U$, and $V$, where*

$$U := \sum_i |i\rangle\langle i| \otimes U_i, \quad V|0^m\rangle = \frac{1}{\sqrt{\alpha}} \sum_i \sqrt{\alpha_i}|i\rangle, \quad \text{and} \quad \alpha := \sum_i \alpha_i, \tag{17}$$

*and outputs a bit indicating whether it was successful.*

*Proof.* By Lemma 7, we can exactly prepare the state $\sum_i \alpha_i T_i(A)|b\rangle/\|\sum_i \alpha_i T_i(A)|b\rangle\|$ with constant success probability and with the stated resource requirements. Since the functions $f$ and $\sum_i \alpha_i T_i$ are $\epsilon$-close on a domain that includes the spectrum of $A$, we have $\|f(A) - \sum_i \alpha_i T_i(A)\| < \epsilon$. Since $|f(x)| \geq 1$ for all $x \in D$, the smallest eigenvalue of $f(A)$ in absolute value is at least 1 and hence by Proposition 9, the output state is $4\epsilon$-close to $f(A)|b\rangle/\|f(A)|b\rangle\|$, as claimed. □

In the following two sections, we apply this lemma to the function $f(x) = 1/x$ on the domain $D_\kappa := [-1, -1/\kappa] \cup [1/\kappa, 1]$, on which $f(x)$ satisfies $|f(x)| \geq 1$. Then we have $\|f(A)|b\rangle\|^{-1} \leq 1$, so the number of uses of $\mathcal{P}_B$, $U$, and $V$ is $O(\alpha)$.

## 3 Fourier approach

We now describe the Fourier approach, which is based on an approximation of $1/A$ as a linear combination of unitaries $e^{-iAt_i}$, $t_i \in \mathbb{R}$. These unitaries can be implemented using any Hamiltonian simulation method. For sparse $A$, we use the method of [BCK15]. Our quantum algorithm for the QLSP then uses Corollary 10 to prepare a quantum state that is $\epsilon$-close to $|x\rangle$. To that end, we establish the following Fourier expansion of the function $1/x$ on the domain $D_\kappa$.

**Lemma 11.** *Let the function $h(x)$ be defined as*

$$h(x) := \frac{i}{\sqrt{2\pi}} \sum_{j=0}^{J-1} \delta_y \sum_{k=-K}^{K} \delta_z \, z_k e^{-z_k^2/2} e^{-ixy_j z_k}, \tag{18}$$

*where $y_j := j\delta_y$, $z_k := k\delta_z$, for some $J = \Theta(\frac{\kappa}{\epsilon}\log(\kappa/\epsilon))$, $K = \Theta(\kappa \log(\kappa/\epsilon))$, $\delta_y = \Theta(\epsilon/\sqrt{\log(\kappa/\epsilon)})$ and $\delta_z = \Theta((\kappa\sqrt{\log(\kappa/\epsilon)})^{-1})$. Then $h(x)$ is is $\epsilon$-close to $1/x$ on the domain $D_\kappa$.*

We start by showing how to represent $1/x$ as a weighted double integral of $e^{-ixt}$ and then approximate the integrals by finite sums. Given any odd function $f : \mathbb{R} \to \mathbb{R}$ satisfying $\int_0^\infty \mathrm{d}y \, f(y) = 1$, we have $1/x = \int_0^\infty \mathrm{d}y \, f(xy)$ for $x \neq 0$. To achieve good performance, we would like both $f$ and its Fourier transform to decay rapidly. Here we choose $f(y) = ye^{-y^2/2}$ (although other choices are possible). This function satisfies

$$f(y) = \frac{i}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathrm{d}z \, z e^{-z^2/2} e^{-iyz} \tag{19}$$

(i.e., it is an eigenfunction of the Fourier transform of eigenvalue $-i$), so

$$\frac{1}{x} = \frac{i}{\sqrt{2\pi}} \int_0^\infty \mathrm{d}y \int_{-\infty}^{\infty} \mathrm{d}z \, z e^{-z^2/2} e^{-ixyz}. \tag{20}$$

We can approximate (20) by a finite sum by restricting to a finite range of integration and discretizing the integral. In the remainder of this section, we analyze this approximation and thereby establish Theorem 3. We consider query complexity in Section 3.1 and then analyze gate complexity in Section 3.2.

### 3.1 Query complexity

The query complexity of this approach is determined by the query complexity of $U$ in Corollary 10 (since the operation $V$ requires no queries to implement). In turn, the query complexity of $U$ depends on the simulation precision and evolution times that are required to obtain an $\epsilon$-approximation of $1/A$. To determine this, it suffices to understand the error introduced by truncating (20) to a finite range of integration.

**Lemma 12.** *The function*

$$g(x) := \frac{i}{\sqrt{2\pi}} \int_0^{y_J} \mathrm{d}y \int_{-z_K}^{z_K} \mathrm{d}z \, z e^{-z^2/2} e^{-ixyz} \tag{21}$$

*is $\epsilon$-close to $1/x$ on the domain $D_\kappa$ for some $y_J = \Theta(\kappa\sqrt{\log(\kappa/\epsilon)})$ and $z_K = \Theta(\sqrt{\log(\kappa/\epsilon)})$.*

*Proof.* Performing the integral over $y$ first, which does not change the value by Fubini's theorem, we have

$$g(x) = \frac{1}{\sqrt{2\pi}x} \int_{-z_K}^{z_K} \mathrm{d}z \, e^{-z^2/2} (1 - e^{-ixy_J z}). \tag{22}$$

Therefore

$$\begin{aligned}
\left| g(x) - \frac{1}{x} \right| &= \left| g(x) - \frac{1}{\sqrt{2\pi}x} \int_{-\infty}^{\infty} \mathrm{d}z \, e^{-z^2/2} \right| && (23)\\
&\leq \frac{1}{\sqrt{2\pi}|x|} \left| \int_{-\infty}^{\infty} \mathrm{d}z \, e^{-z^2/2} e^{-ixy_J z} \right| + \frac{4}{\sqrt{2\pi}|x|} \int_{z_K}^{\infty} \mathrm{d}z \, e^{-z^2/2} && (24)\\
&= \frac{1}{|x|} e^{-(xy_J)^2/2} + \frac{4}{\sqrt{2\pi}|x|} \int_{z_K}^{\infty} \mathrm{d}z \, e^{-z^2/2} && (25)\\
&\leq \frac{1}{|x|} e^{-(xy_J)^2/2} + \frac{2}{|x|} e^{-z_K^2/2}, && (26)
\end{aligned}$$

where we used the bound

$$\frac{1}{\sqrt{2\pi}} \int_{z_K}^{\infty} \mathrm{d}z \, e^{-z^2/2} \leq \frac{1}{2} e^{-z_K^2/2} \tag{27}$$

in the last step. Since $|x| \geq 1/\kappa$, there exist $y_J = \Theta(\kappa\sqrt{\log(\kappa/\epsilon)})$ and $z_K = \Theta(\sqrt{\log(\kappa/\epsilon)})$ such that (26) is at most $\epsilon$. □

To apply Corollary 10, we discretize this integral, approximating $g(x)$ by the function $h(x)$ defined in Lemma 11. In particular, $y_J$ and $z_K$ are as in the definition of $g(x)$ in Lemma 12. By taking $\delta_y$ and $\delta_z$ sufficiently small, $h(x)$ can approximate $g(x)$ arbitrarily closely. Since the query complexity resulting from Corollary 10 does not depend on the number of terms in the linear combination, we can make the discretization error arbitrarily small and neglect its contribution. This shows the query complexity part of Theorem 3, as follows.

*Proof of Theorem 3 (query complexity).* We implement $h(A)$ using Corollary 10. The $L_1$ norm of the coefficients of this linear combination is

$$\alpha = \frac{1}{\sqrt{2\pi}} \sum_{j=0}^{J-1} \delta_y \sum_{k=-K}^{K} \delta_z \, |z_k| e^{-z_k^2/2} = \Theta(y_J), \tag{28}$$

where we used the fact that

$$\sum_{k=-K}^{K} \delta_z \, |z_k| e^{-z_k^2/2} \approx \int_{-z_K}^{z_K} \mathrm{d}z \, |z| e^{-z^2/2} \leq 2 \int_0^{\infty} \mathrm{d}z \, z e^{-z^2/2} = 2 \tag{29}$$

(here the error in the approximation is negligible since we can take $\delta_z$ arbitrarily small without affecting the query complexity). By Lemma 12, we have $\alpha = O(\kappa\sqrt{\log(\kappa/\epsilon)})$. The evolution times of the Hamiltonian simulations that appear in the linear combination are $t = O(y_J z_K) = O(\kappa \log(\kappa/\epsilon))$. Since we invoke Hamiltonian simulation $\Theta(\alpha)$ times, the overall error is at most $\epsilon$ provided each simulation has error $\epsilon' = O(\epsilon/\alpha) = O(\epsilon/(\kappa\sqrt{\log(\kappa/\epsilon)}))$.

The overall query complexity is the number of uses of Hamiltonian simulation times the query complexity of each simulation. Using the Hamiltonian simulation algorithm in [BCK15, Lemma 9], the query complexity of simulating a $d$-sparse Hamiltonian for time $t$ with error at most $\epsilon'$ is $O(d\|A\|_{\max} t \log(\|A\|t/\epsilon')/\log\log(\|A\|t/\epsilon')) = O(dt\log(t/\epsilon'))$, since for the QLSP, we have $\|A\|_{\max} \leq \|A\| \leq 1$. Thus the total query complexity is

$$O(\alpha dt \log(t/\epsilon')) = O\Big(\kappa\sqrt{\log(\kappa/\epsilon)} d\kappa \log(\kappa/\epsilon) \log\Big(\frac{\kappa^2}{\epsilon} \log^{1.5}(\kappa/\epsilon)\Big)\Big) = O(d\kappa^2 \log^{2.5}(\kappa/\epsilon)). \tag{30}$$

Finally, by Corollary 10, the number of uses of $\mathcal{P}_B$ is $O(\alpha) = O(\kappa\sqrt{\log(\kappa/\epsilon)})$. □

### 3.2 Gate complexity

The gate complexity of this approach is given by the gate complexity of the unitaries in Corollary 10, namely quantum simulation and the state preparation map $V$, times the number of amplitude amplification steps. The cutoffs $J$ and $K$ (which determine $\delta_y$ and $\delta_z$) affect the gate complexity as they determine the number of coefficients appearing in the approximation of $1/A$ by $h(A)$.

Our analysis uses the following identity.

**Lemma 13.** *Let $\omega \in \mathbb{R}$ and $\delta_z > 0$. Then*

$$\sum_{k=-\infty}^{\infty} e^{-(\omega + 2\pi k/\delta_z)^2/2} = \frac{1}{\sqrt{2\pi}} \sum_{k=-\infty}^{\infty} \delta_z \, e^{-z_k^2/2} e^{-i\omega z_k} \tag{31}$$

*where $z_k := k\delta_z$.*

*Proof.* Equation (31) is a case of the Poisson summation formula, which states that if $f\colon \mathbb{R} \to \mathbb{C}$ is a Schwartz function with Fourier transform $\hat{f}\colon \mathbb{R} \to \mathbb{C}$, then

$$\sum_{k=-\infty}^{\infty} f(k) = \sqrt{2\pi} \sum_{k=-\infty}^{\infty} \hat{f}(2\pi k). \tag{32}$$

Here $f$ is called a Schwartz function if, for all nonnegative integers $m$ and $n$, $\sup_{x\in\mathbb{R}} |x^m \frac{\mathrm{d}^n}{\mathrm{d}x^n} f(x)|$ is at most some constant (that can depend on $m$ and $n$).

In our case, $f(x) = e^{-(\omega + 2\pi x/\delta_z)^2/2}$, with the Fourier transform $\hat{f}(y) = \frac{\delta_z}{2\pi} e^{-(y\delta_z)^2/8\pi^2} e^{-iy\delta_z\omega/2\pi}$. Then the left-hand side of (32) coincides with the left-hand side of (31), and it is easy to check that the right-hand side of (32) coincides with the right-hand side of (31). It is well known that a Gaussian is a Schwartz function, so the identity follows. □

Now we quantify how well $h(x)$ in (18) approximates $1/x$.

*Proof of Lemma 11.* We have $\delta_y = y_J/J = \Theta(\epsilon/\sqrt{\log(\kappa/\epsilon)})$ and $\delta_z = z_K/K = \Theta(1/\kappa\sqrt{\log(\kappa/\epsilon)})$.

Performing the geometric sum over $j$, we have

$$h(x) = \frac{i\delta_y}{\sqrt{2\pi}} \sum_{k=-K}^{K} \delta_z \, z_k e^{-z_k^2/2} \frac{1 - e^{-ixy_J z_k}}{1 - e^{-ix\delta_y z_k}}. \tag{33}$$

Using the triangle inequality, we have the bound

$$\left|h(x) - \frac{1}{x}\right| \le \left| \frac{i\delta_y}{\sqrt{2\pi}} \sum_{k=-K}^{K} \delta_z \, z_k e^{-z_k^2/2} \frac{1 - e^{-ixy_J z_k}}{1 - e^{-ix\delta_y z_k}} - \frac{1}{\sqrt{2\pi}x} \sum_{k=-K}^{K} \delta_z \, e^{-z_k^2/2}(1 - e^{-ixy_J z_k}) \right| \tag{34}$$

$$+ \left| \frac{1}{\sqrt{2\pi}x} \sum_{k=-K}^{K} \delta_z \, e^{-z_k^2/2}(1 - e^{-ixy_J z_k}) - \frac{1}{\sqrt{2\pi}x} \sum_{k=-\infty}^{\infty} \delta_z \, e^{-z_k^2/2}(1 - e^{-ixy_J z_k}) \right| \tag{35}$$

$$+ \left| \frac{1}{\sqrt{2\pi}x} \sum_{k=-\infty}^{\infty} \delta_z \, e^{-z_k^2/2}(1 - e^{-ixy_J z_k}) - \frac{1}{\sqrt{2\pi}x} \sum_{k=-\infty}^{\infty} \delta_z \, e^{-z_k^2/2} \right| \tag{36}$$

$$+ \left| \frac{1}{\sqrt{2\pi}x} \sum_{k=-\infty}^{\infty} \delta_z \, e^{-z_k^2/2} - \frac{1}{x} \right|. \tag{37}$$

We show that each term on the right-hand side is $O(\epsilon)$.

Since $|\frac{1}{1-e^{-ix}} - \frac{1}{ix}| < 1$ for all $x \in [-1, 1]$ and since $|x\delta_y z_K| = O(\epsilon) < 1$ for sufficiently small $\epsilon$, the term on the right-hand side of (34) is upper bounded by

$$\frac{\delta_y}{\sqrt{2\pi}} \sum_{k=-K}^{K} \delta_z \, |z_k| e^{-z_k^2/2} \le \frac{\delta_y}{\sqrt{2\pi}} \sum_{k=-K}^{K} \delta_z \, e^{-z_k^2/4} \tag{38}$$

$$\le \sqrt{\frac{2}{\pi}} \delta_y \int_0^{\infty} \mathrm{d}z \, e^{-z^2/4} \tag{39}$$

$$= \sqrt{2}\delta_y = O(\epsilon). \tag{40}$$

Here in (38) we used the fact that $|x|e^{-x^2/2} < e^{-x^2/4}$ for all $x \in \mathbb{R}$.

The term in (35) is

$$\left| \frac{1}{\sqrt{2\pi}x} \sum_{|k|>K} \delta_z \, e^{-z_k^2/2}(1 - e^{-ixy_J z_k}) \right| \le \sqrt{\frac{2}{\pi}} \frac{2}{|x|} \sum_{k=K+1}^{\infty} \delta_z \, e^{-z_k^2/2} \tag{41}$$

$$\le \sqrt{\frac{2}{\pi}} \frac{2}{|x|} \int_{z_K}^{\infty} \mathrm{d}z \, e^{-z^2/2} \tag{42}$$

$$\le \frac{2}{|x|} e^{-z_K^2/2} = O(\epsilon) \tag{43}$$

where we upper bounded the sum of a decreasing sequence by the corresponding integral and applied (27).

Using Lemma 13 with $\omega = xy_J$, the term in (36) is

$$\frac{1}{\sqrt{2\pi}|x|} \sum_{k=-\infty}^{\infty} \delta_z \, e^{-z_k^2/2} e^{-ixy_J z_k} = \frac{1}{|x|} \sum_{k=-\infty}^{\infty} e^{-(xy_J + 2\pi k/\delta_z)^2/2}. \tag{44}$$

For $k \neq 0$, we have $|xy_J + 2\pi k/\delta_z| \geq |k|(2\pi/\delta_z - y_J)$. By choosing $K$ sufficiently large, we can ensure that $2\pi/\delta_z \geq y_J$. Thus we see that (36) is at most

$$\frac{1}{|x|}\left(e^{-(xy_J)^2/2} + 2\sum_{k=1}^{\infty} e^{-k(2\pi/\delta_z - y_J)^2}\right) = \frac{1}{|x|}\left(e^{-(xy_J)^2/2} + \frac{2}{e^{(2\pi/\delta_z - y_J)^2} - 1}\right) = O(\epsilon). \tag{45}$$

Finally, by Lemma 13 with $\omega = 0$, we have

$$\frac{1}{\sqrt{2\pi}} \sum_{k=-\infty}^{\infty} \delta_z \, e^{-z_k^2/2} = \sum_{k=-\infty}^{\infty} e^{-(2\pi k/\delta_z)^2/2} \tag{46}$$

$$\leq 1 + 2\sum_{k=1}^{\infty} e^{-2\pi^2 k/\delta_z^2} \tag{47}$$

$$= 1 + \frac{2}{e^{2\pi^2/\delta_z^2} - 1} \tag{48}$$

$$= 1 + O(\epsilon \, e^{-\kappa^2 \log \kappa}), \tag{49}$$

so (37) is also $O(\epsilon)$, completing the proof. □

We can now determine the gate complexity of the Fourier approach to the QLSP.

*Proof of Theorem 3 (gate complexity).* To apply Corollary 10, we must implement a unitary $V$ that maps the $m$-qubit state $|0^m\rangle$, where $m = O(\log(JK))$, to

$$\frac{1}{(2\pi)^{\frac{1}{4}}\sqrt{\alpha}} \sum_{j=0}^{J-1} \sqrt{\delta_y} \sum_{k=-K}^{K} \sqrt{\delta_z |z_k|} e^{-z_k^2/2} |j,k\rangle = \frac{\sqrt{\delta_y}\sqrt{\delta_z}}{(2\pi)^{\frac{1}{4}}\sqrt{\alpha}} \left(\sum_{j=0}^{J-1} |j\rangle\right) \otimes \left(\sum_{k=-K}^{K} \sqrt{|z_k|} e^{-z_k^2/2} |k\rangle\right). \tag{50}$$

From the query complexity part of the proof of Theorem 3, we have $\alpha = O(\kappa\sqrt{\log(\kappa/\epsilon)})$. The operation $V$ can be implemented in two steps, preparing the superpositions over $|j\rangle$ and $|k\rangle$ independently since this is a product state. First we use $O(\log(J))$ Hadamard gates to prepare $\sum_{j=0}^{J-1} |j\rangle/\sqrt{J}$ (assuming for simplicity that $J$ is a power of 2). Then we use $O(K)$ gates to prepare the corresponding superposition over $|k\rangle$ [SBM06]. Using the values of $J$ and $K$ from Lemma 11, the gate complexity of $V$ is $O(\log(J) + K) = O(\log(\kappa \log(\kappa/\epsilon)/\epsilon) + \kappa \log(\kappa/\epsilon)) = O(\kappa \log(\kappa/\epsilon))$.

Corollary 10 also requires the unitary

$$U = i \sum_{j=0}^{J-1} \sum_{k=-K}^{K} |j,k\rangle\langle j,k| \otimes \operatorname{sgn}(k) e^{-iAy_j z_k}. \tag{51}$$

Using Lemma 8 (specifically, in the form discussed immediately following the proof of the lemma) and the algorithm of [BCK15, Lemma 10], we can simulate $U$ within precision $\epsilon'$ with gate complexity $O(d\|A\|_{\max} t (\log N + \log^{2.5}(\|A\|t/\epsilon')) \log(\|A\|t/\epsilon')/\log\log(\|A\|t/\epsilon'))$, where $t = y_J z_K$. This complexity is $O(dt(\log N + \log^{2.5}(t/\epsilon')) \log(t/\epsilon'))$, since for the QLSP, we have $\|A\|_{\max} \leq \|A\| \leq 1$. Since $y_J z_K = O(\kappa \log(\kappa/\epsilon))$ by Lemma 12 and $\epsilon' \leq \epsilon$, the operation $U$ has gate complexity $O(d\kappa \log^2(\kappa/\epsilon')(\log N + \log^{2.5}(\kappa/\epsilon')))$. Note that the gate complexity of $V$ is dominated by that of $U$.

Since we invoke Hamiltonian simulation $O(\alpha)$ times, with $\alpha = O(\kappa\sqrt{\log(\kappa/\epsilon)})$ (see the proof of the query complexity part of Theorem 3), we can choose $\epsilon' = O(\epsilon/\alpha)$ for overall error at most $\epsilon$. Thus the overall gate complexity of the method is

$$O(d\kappa^2 \log^{2.5}(\kappa/\epsilon)(\log N + \log^{2.5}(\kappa/\epsilon))) \tag{52}$$

as claimed. □

# 4 Chebyshev approach

We now describe our second approach to the QLSP. This approach uses a Chebyshev expansion to implement $A^{-1}$ without appealing directly to Hamiltonian simulation. Instead of building the function $1/x$ as a linear combination of terms of the form $e^{-itx}$ (as in Section 3), we approximate it by a linear combination of terms of the form $\mathcal{T}_n(x)$, where $\mathcal{T}_n$ is the $n^{\text{th}}$ Chebyshev polynomial of the first kind. We implement such terms using a quantum walk that has previously been applied to Hamiltonian simulation [Chi10, BC12, BCK15]. With that goal, we establish the following decomposition of $1/x$ as a linear combination of Chebyshev polynomials.

**Lemma 14.** *Let $g(x)$ be defined as*

$$g(x) := 4\sum_{j=0}^{j_0}(-1)^j \left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right] \mathcal{T}_{2j+1}(x), \tag{53}$$

*where $j_0 = 2\sqrt{b\log(4b/\epsilon)}$ and $b = \kappa^2\log(\kappa/\epsilon)$. Then $g(x)$ is $2\epsilon$-close to $1/x$ on $D_\kappa$.*

This section is organized as follows. In Section 4.1 we define a quantum walk for any given Hamiltonian and express this walk (and its powers) in terms of Chebyshev polynomials. In Section 4.2, we present an expansion of $1/x$ in terms of Chebyshev polynomials and use this to establish the query complexity of the Chebyshev approach to the QLSP. In Section 4.3, we upper bound the gate complexity of this approach.

## 4.1 A quantum walk for any Hamiltonian

Let $A$ be a $d$-sparse $N \times N$ Hamiltonian with $\|A\|_{\max} \leq 1$. We now define a quantum walk corresponding to $A$ and express its action in terms of Chebyshev polynomials.

The quantum walk is defined using a set of states $\{|\psi_j\rangle \in \mathbb{C}^{2N} \otimes \mathbb{C}^{2N} : j \in [N]\}$ defined as

$$|\psi_j\rangle := |j\rangle \otimes \frac{1}{\sqrt{d}} \sum_{k\in[N]} \left(\sqrt{A_{jk}^*}|k\rangle + \sqrt{1-|A_{jk}|}\,|k+N\rangle\right). \tag{54}$$

Note that the square root in (54) is potentially ambiguous when $A_{jk}$ is complex. Our results hold for any consistent choice of square root that ensures $\sqrt{A_{jk}^*}\left(\sqrt{A_{jk}}\right)^* = A_{jk}^*$. For more detail, see the discussion preceding [BC12, eq. 14].

The quantum walk now occurs in the Hilbert space $\mathbb{C}^{2N} \otimes \mathbb{C}^{2N}$. We define an isometry $T$ from $\mathbb{C}^N$ to the walk space by

$$T := \sum_{j\in[N]} |\psi_j\rangle\langle j|. \tag{55}$$

We define a swap operator on the walk space by $S|j,k\rangle = |k,j\rangle$ for all $j,k \in [2N]$. Finally, the walk operator is $W := S(2TT^\dagger - \mathbb{1})$. As discussed in [BC12, BCK15], the walk operator $W$ can be implemented using $O(1)$ queries to $\mathcal{P}_A$, and this implementation is gate-efficient as we discuss in Section 4.3.

We now analyze the structure of this walk.

**Lemma 15.** *Within the invariant subspace* $\mathrm{span}\{T|j\rangle, ST|j\rangle : j \in [N]\}$*, the walk operator* $W$ *has the block form*

$$\begin{pmatrix} H & -\sqrt{1-H^2} \\ \sqrt{1-H^2} & H \end{pmatrix}$$

*where* $H := A/d$ *and where the first block consists of the states* $\{T|j\rangle : j \in [N]\}$*.*

*Proof.* Observe that $T^\dagger T = \mathbb{1}$ and $T^\dagger S T = H$. Let $|\lambda\rangle$ denote an eigenvector of $H$ with eigenvalue $\lambda \in [-1,1]$. Let $|\perp_\lambda\rangle$ denote a state orthogonal to $T|\lambda\rangle$ in $\mathrm{span}\{T|\lambda\rangle, ST|\lambda\rangle\}$, satisfying

$$ST|\lambda\rangle = \lambda T|\lambda\rangle + \sqrt{1-\lambda^2}|\perp_\lambda\rangle. \tag{56}$$

Then we have

$$\begin{aligned} WT|\lambda\rangle &= S(2TT^\dagger - 1)T|\lambda\rangle && (57)\\ &= ST|\lambda\rangle && (58)\\ &= \lambda T|\lambda\rangle + \sqrt{1-\lambda^2}|\perp_\lambda\rangle. && (59) \end{aligned}$$

Furthermore,

$$\begin{aligned} \sqrt{1-\lambda^2}W|\perp_\lambda\rangle &= WST|\lambda\rangle - \lambda WT|\lambda\rangle && (60)\\ &= S(2TT^\dagger - 1)ST|\lambda\rangle - \lambda S(2TT^\dagger - 1)T|\lambda\rangle && (61)\\ &= \lambda ST|\lambda\rangle - T|\lambda\rangle && (62)\\ &= (\lambda^2 - 1)T|\lambda\rangle + \lambda\sqrt{1-\lambda^2}|\perp_\lambda\rangle, && (63) \end{aligned}$$

so

$$W|\perp_\lambda\rangle = -\sqrt{1-\lambda^2}T|\lambda\rangle + \lambda|\perp_\lambda\rangle. \tag{64}$$

Combining (59) and (64), we see that $W$ has the form

$$\begin{pmatrix} \lambda & -\sqrt{1-\lambda^2} \\ \sqrt{1-\lambda^2} & \lambda \end{pmatrix} \tag{65}$$

within the invariant subspace $\mathrm{span}\{T|\lambda\rangle, |\perp_\lambda\rangle\}$. Since $\mathrm{span}\{|j\rangle : j \in [N]\} = \mathrm{span}\{|\lambda\rangle\}$, the result follows. □

**Lemma 16.** *Let*

$$W = \begin{pmatrix} \lambda & -\sqrt{1-\lambda^2} \\ \sqrt{1-\lambda^2} & \lambda \end{pmatrix}.$$

*Then for any positive integer* $n$*,*

$$W^n = \begin{pmatrix} \mathcal{T}_n(\lambda) & -\sqrt{1-\lambda^2}\,\mathcal{U}_{n-1}(\lambda) \\ \sqrt{1-\lambda^2}\,\mathcal{U}_{n-1}(\lambda) & \mathcal{T}_n(\lambda) \end{pmatrix}$$

*where* $\mathcal{T}_n$ *is the* $n^{th}$ *Chebyshev polynomial of the first kind and* $\mathcal{U}_n$ *is the* $n^{th}$ *Chebyshev polynomial of the second kind.*

*Proof.* This follows by straightforward induction. The base case ($n = 1$) follows from the identities $\mathcal{T}_1(\lambda) = \lambda$ and $\mathcal{U}_0(\lambda) = 1$. Assuming the claim holds for a given value of $n$, we have

$$W^n W = \begin{pmatrix} \lambda\,\mathcal{T}_n(\lambda) - (1-\lambda^2)\,\mathcal{U}_{n-1}(\lambda) & -\sqrt{1-\lambda^2}\,(\mathcal{T}_n(\lambda) + \lambda\,\mathcal{U}_{n-1}(\lambda)) \\ \sqrt{1-\lambda^2}\,(\lambda\,\mathcal{U}_{n-1}(\lambda) + \mathcal{T}_n(\lambda)) & -(1-\lambda^2)\,\mathcal{U}_{n-1}(\lambda) + \lambda\,\mathcal{T}_n(\lambda) \end{pmatrix} \tag{66}$$

$$= \begin{pmatrix} \mathcal{T}_{n+1}(\lambda) & -\sqrt{1-\lambda^2}\,\mathcal{U}_n(\lambda) \\ \sqrt{1-\lambda^2}\,\mathcal{U}_n(\lambda) & \mathcal{T}_{n+1}(\lambda) \end{pmatrix} \tag{67}$$

where we used the identities

$$\mathcal{T}_{n+1}(\lambda) = \lambda\,\mathcal{T}_n(\lambda) - (1-\lambda^2)\,\mathcal{U}_{n-1}(\lambda) \tag{68}$$

$$\mathcal{U}_n(\lambda) = \mathcal{T}_n(\lambda) + \lambda\,\mathcal{U}_{n-1}(\lambda) \tag{69}$$

relating Chebyshev polynomials of the first and second kinds. □

Therefore, using $n$ steps of this walk, we can implement a transformation of the form

$$|0^m\rangle|\psi\rangle \mapsto |0^m\rangle\mathcal{T}_n(H)|\psi\rangle + |\Phi^\perp\rangle, \tag{70}$$

where $|\Phi^\perp\rangle$ is an unnormalized state satisfying $\Pi|\Phi^\perp\rangle = 0$, where $\Pi := |0^m\rangle\langle 0^m| \otimes \mathbb{1}$.

Finally, we note that since the walk operator $W$ can be implemented using $O(1)$ queries to the oracle $\mathcal{P}_A$, as described in [BC12, BCK15], the operation in (70) can be performed using $O(n)$ queries.

### 4.2 Query complexity

The Hamiltonian simulation algorithm of [BCK15] approximates $e^{-iHt}$ by a linear combination of the matrices $W^n$, where $W$ is the quantum walk defined above. From the form of $W^n$ in terms of Chebyshev polynomials, we can interpret the algorithm as decomposing $e^{-iHt}$ as a linear combination of the Chebyshev polynomials $\mathcal{T}_n(H)$.

Applying a similar approach to linear systems, we aim to approximate $H^{-1}$, where $H = A/d$, by a linear combination of $\mathcal{T}_n(H)$. Since $H$ is Hermitian, it suffices to find a good approximation of $f(x) = 1/x$ by a linear combination of $\mathcal{T}_n(x)$ over a domain containing the spectrum of $H$. Since the eigenvalues of $A$ lie in $D_\kappa := [-1, -1/\kappa] \cup [1/\kappa, 1]$, the eigenvalues of $H$ lie in $D_{\kappa d}$. (Although the eigenvalues of $H$ actually lie in $[-1/d, -1/(\kappa d)] \cup [1/(\kappa d), 1/d]$, we do not use this fact.) Once we express $f(x) = 1/x$ as a linear combination of $\mathcal{T}_n(x)$ over the domain $D_{\kappa d}$, the same linear combination yields an approximation for $H^{-1}$, according to Corollary 10.

We now analyze the approximation of $1/x$ by a Chebyshev series. We start by approximating the function $1/x$ with a function that is bounded at the origin. We accomplish this by multiplying $1/x$ with a function that is close to identity on $D_{\kappa d}$ and very small near the origin. The function $1 - (1 - x^2)^b$ has this property for large enough $b$.

**Lemma 17.** *The function*

$$f(x) := \frac{1 - (1-x^2)^b}{x} \tag{71}$$

*is $\epsilon$-close to $1/x$ on the domain $D_{\kappa d}$ for any integer $b \geq (\kappa d)^2 \log(\kappa d/\epsilon)$.*

*Proof.* For $b > 0$, on the domain $D_{\kappa d}$ the numerator $1-(1-x^2)^b$ increases monotonically toward 1 as $|x|$ increases. Thus, over $D_{\kappa d}$, the numerator differs most from the constant function 1 at $x = 1/(\kappa d)$, where the difference is

$$\left(1 - \frac{1}{(\kappa d)^2}\right)^b \leq e^{-b/(\kappa d)^2} = \frac{\epsilon}{\kappa d}. \tag{72}$$

Therefore we have

$$\left|\frac{1-(1-x^2)^b}{x} - \frac{1}{x}\right| \leq \frac{\epsilon}{\kappa d|x|} \leq \epsilon. \qquad \square \tag{73}$$

We now express the function $f(x)$ as a linear combination of Chebyshev polynomials. Since $f(x)$ is a polynomial of degree $2b-1$, it can be exactly represented as a linear combination of Chebyshev polynomials of order at most $2b-1$.

**Lemma 18.** *Over the domain $[-1,1]$, the function $f(x)$ defined in (71) can be exactly represented by a linear combination of Chebyshev polynomials of order at most $2b-1$:*

$$f(x) = \frac{1-(1-x^2)^b}{x} = 4\sum_{j=0}^{b-1}(-1)^j \left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right] \mathcal{T}_{2j+1}(x). \tag{74}$$

*Proof.* First note that $f(x)$ is well-defined as $x \to 0$ because its numerator is a polynomial with no constant term.

Since $x \in [-1,1]$, we can make the substitution $x = \cos(\theta)$. Thus we aim to prove

$$f(\cos(\theta)) = \frac{1-(1-\cos^2(\theta))^b}{\cos(\theta)} = 4\sum_{j=0}^{b-1}(-1)^j \left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right] \mathcal{T}_{2j+1}(\cos(\theta)). \tag{75}$$

Using $\mathcal{T}_n(\cos\theta) = \cos(n\theta)$, this is equivalent to showing

$$1-(\sin(\theta))^{2b} = 4\sum_{j=0}^{b-1}(-1)^j \left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right] \cos((2j+1)\theta)\cos(\theta). \tag{76}$$

By the standard power reduction formula for even powers of $\sin\theta$, the left-hand side of (76) simplifies to

$$1-(\sin(\theta))^{2b} = 1 - \left(\frac{1}{2^{2b}}\binom{2b}{b} + \frac{2}{2^{2b}}\sum_{k=0}^{b-1}(-1)^{b-k}\binom{2b}{k}\cos((2b-2k)\theta)\right). \tag{77}$$

Using the product-to-sum formula for cosines, the right-hand side of (76) simplifies to

$$2\sum_{j=0}^{b-1}(-1)^j \left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right] (\cos((2j+2)\theta) + \cos(2j\theta)) \tag{78}$$

It remains to check that the coefficients of $\cos(2r\theta)$ are equal on both sides for all $r \in \{0,1,\ldots,b\}$. (Note that both sides are linear combinations of such terms.)

First, the coefficient of the $r=0$ term on the left-hand side is $1 - \frac{1}{2^{2b}}\binom{2b}{b}$. The same coefficient on the right-hand side is $2^{1-2b}\sum_{i=1}^{b}\binom{2b}{b+i} = 2^{-2b}\sum_{i=0}^{2b}\binom{2b}{i} - \binom{2b}{b} = 1 - \frac{1}{2^{2b}}\binom{2b}{b}$.

When $r \in [b]$, the coefficient of $\cos(2r\theta)$ on the left-hand side is

$$-\left(\frac{2}{2^{2b}}(-1)^r\binom{2b}{b-r}\right). \tag{79}$$

When $r \in [b-1]$, the coefficient of $\cos(2r\theta)$ on the right-hand side is

$$2(-1)^{r-1}\left[\frac{\sum_{i=r}^{b}\binom{2b}{b+i}}{2^{2b}}\right] - 2(-1)^r\left[\frac{\sum_{i=r+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right] = 2(-1)^{r-1}\left[\frac{\binom{2b}{b+r}}{2^{2b}}\right], \tag{80}$$

which is the same as the left-hand side.

Finally, when $r = b$, the coefficient of $\cos(2r\theta)$ on the right-hand side is

$$2(-1)^{b-1}\left[\frac{\sum_{i=b}^{b}\binom{2b}{b+i}}{2^{2b}}\right] = 2(-1)^{b-1}\left[\frac{1}{2^{2b}}\right], \tag{81}$$

which is also the same as the left-hand side. □

Although $f(x)$ is exactly representable as a linear combination of Chebyshev polynomials of order at most $2b-1$, observe that the coefficients corresponding to the higher order terms are very small. Thus we can truncate the series expansion for $f(x)$ while still remaining $\epsilon$-close to $f(x)$.

**Lemma 19.** *The function $f(x)$ in (71) can be $\epsilon$-approximated by a linear combination of Chebyshev polynomials of order $O(\sqrt{b\log(b/\epsilon)})$ by truncating the series in (74) at $j_0 = 2\sqrt{b\log(4b/\epsilon)}$. In other words,*

$$g(x) := 4\sum_{j=0}^{j_0}(-1)^j\left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right]\mathcal{T}_{2j+1}(x) \tag{82}$$

*is $\epsilon$-close to $f(x)$ on $[-1,1]$.*

*Proof.* The expression in square brackets in (74) is the probability of seeing more than $b+j$ heads on flipping $2b$ fair coins, which is negligible for large $j$. In particular, the Chernoff bound gives

$$\frac{1}{2^{2b}}\sum_{i=j+1}^{b}\binom{2b}{b+i} \le e^{-j^2/4b}. \tag{83}$$

Thus we have

$$|f(x)-g(x)| = 4\left|\sum_{j=j_0+1}^{b-1}(-1)^j\left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right]\mathcal{T}_{2j+1}(x)\right| \tag{84}$$

$$\le 4\sum_{j=j_0+1}^{b-1}e^{-j^2/4b}|\mathcal{T}_{2j+1}(x)| \quad \le 4be^{-j_0^2/4b} \quad = \epsilon, \tag{85}$$

where we used the fact that $|\mathcal{T}_n(x)| \le 1$ for $x \in [-1,1]$. □

We now establish the query complexity part of Theorem 4.

*Proof of Theorem 4 (query complexity).* We implement $A^{-1}$ using Corollary 10 by expressing it in terms of $g(H)$. Since $\|H^{-1} - g(H)\| \leq \epsilon$, $\|A^{-1} - \frac{1}{d}g(H)\| \leq \epsilon/d \leq \epsilon$. Corollary 10 makes $O(\alpha)$ uses of $U$, $V$, and $\mathcal{P}_B$. The value of $\alpha$ for the linear combination $\frac{1}{d}g(H)$ is

$$\alpha = \frac{4}{d}\sum_{j=0}^{j_0}\left[\frac{\sum_{i=j+1}^{b}\binom{2b}{b+i}}{2^{2b}}\right] \leq \frac{4j_0}{d} \tag{86}$$

because the term in square brackets is a probability and hence is at most 1. Since $V$ costs no queries to implement, it suffices to understand the cost of implementing $U$.

The highest order of a Chebyshev polynomial used in (82) is $O(j_0) = O(d\kappa \log(d\kappa/\epsilon))$. Since the cost of implementing the map (70) is proportional to $n$ as discussed in Section 4.1, the cost of implementing the unitary $U$ in Lemma 7 is $O(j_0)$. Thus the total query complexity is $O(\alpha j_0) = O(d\kappa^2 \log^2(d\kappa/\epsilon))$ Finally, the total number of applications of $\mathcal{P}_B$ is $O(\alpha) = O(\kappa \log(\kappa d/\epsilon))$. □

### 4.3 Gate complexity

To upper bound the gate complexity of this approach, we must consider the implementation of the operator $U$ (which consists of steps of the walk $W$) and the operator $V$ in Corollary 10.

*Proof of Theorem 4 (gate complexity).* We begin with the gate cost of $W$. As discussed in the proof of [BCK15, Lemma 10], a step of the quantum walk can be performed with gate complexity $O(\log N + \log^{2.5}(\kappa d/\epsilon))$. Thus the gate cost of implementing the map $U$, which involves a power of $W$ as large as $j_0$, is $O(j_0(\log N + \log^{2.5}(\kappa d/\epsilon)))$ using Lemma 8.

Implementing the map $V$ involves creating a quantum state in a space of dimension $j_0$. As shown in [SBM06], any such state can be created with $O(j_0)$ two-qubit gates. Since the gate cost of $V$ is less than the gate cost of $U$, we can neglect this in our calculations.

Thus we see that the gate complexity exceeds the query complexity by a multiplicative factor of $O(\log N + \log^{2.5}(\kappa d/\epsilon))$, as claimed. □

## 5 Improved dependence on condition number

In this section we reduce the $\kappa$-dependence of both our algorithms from quadratic to nearly linear, establishing Theorem 5. We use a tool called variable-time amplitude amplification, a low-precision version of phase estimation, and results from Section 3 and Section 4.

To see why this improvement might even be possible, observe that the quadratic dependence on $\kappa$ has two contributions. First, the complexity of performing the linear combination corresponding to $A^{-1}$ (even with very low probability) depends linearly on $\kappa$. More precisely, this refers to the cost of implementing the controlled unitary $U$ in Corollary 10: in the Fourier approach this cost depends on the length of time for which we simulate a Hamiltonian, and in the Chebyshev approach it depends on the highest-order Chebyshev polynomial used. In both approaches, this cost is close to linear in $\kappa$. The second contribution comes from amplifying the probability of $A^{-1}$ being successfully applied. As described in Corollary 10, this contribution is proportional to the quantity $\alpha$, whose dependence on $\kappa$ is also approximately linear in both approaches.

Observe that if we were promised that the state $|b\rangle$ were only a superposition of eigenvectors of $A$ with singular values that are all close to 1, the problem would be much easier. As noted in Section 1.2, if we know that $|b\rangle$ lives in a subspace of $A$ with low condition number $\kappa'$, we can replace $\kappa$ with $\kappa'$ in our upper bounds. If $|b\rangle$ lives in the subspace of eigenvectors of $A$ with singular values close to 1, then $\kappa' = O(1)$. This suggests that the difficult case is when the

state $|b\rangle$ is a superposition of eigenvectors with singular values close to $1/\kappa$. However, if we were promised that this holds, then we could again improve the complexity of our algorithms. This is because the second contribution in Corollary 10, proportional to $\alpha$, only arises because we need to amplify the success probability of our implementation of $A^{-1}$. If $|b\rangle$ were simply a superposition of eigenvectors with singular value close to $1/\kappa$, then this step would be considerably less expensive. Quantitatively, observe that when we apply Corollary 10, the complexity is $O(\alpha/\||A^{-1}|b\rangle\|)$, which would be $O(\alpha/\kappa) = \mathrm{poly}(\log(d\kappa/\epsilon))$ in this case, bringing down the cost by a factor of $\kappa$. Thus at both extremes we can improve the $\kappa$-dependence of our algorithm. The reason our algorithm has nearly quadratic dependence on $\kappa$ is that we need to handle both cases simultaneously.

In the HHL algorithm, a similar situation also arises. Ambainis exploited this to reduce the complexity of the HHL algorithm using a technique called variable-time amplitude amplification (VTAA) [Amb12], a generalization of amplitude amplification. VTAA amplifies the success probability of a quantum algorithm $\mathcal{A}$ with variable stopping times using a sequence of concatenated amplitude amplification steps. If some branches of the computation performed by $\mathcal{A}$ run longer than others, VTAA can perform significantly better than standard amplitude amplification.

### 5.1 Tools

We now describe the tools required to construct our algorithm: variable-time amplitude amplification (Theorem 21), a low-precision variant of phase estimation we call gapped phase estimation (Lemma 22), and procedures for implementing approximations of $A^{-1}$ that are accurate for different ranges of eigenvalues (Lemma 23). In this section we use the notation $\widetilde{O}(\cdot)$ to ignore factors of $\log(d\kappa/\epsilon)$, i.e., we write $f(d,\kappa,\epsilon) = \widetilde{O}(g(d,\kappa,\epsilon))$ to indicate $f(d,\kappa,\epsilon) = O(g(d,\kappa,\epsilon)\,\mathrm{poly}(\log(d\kappa/\epsilon)))$.

**Variable-time amplitude amplification** We start by formalizing the notion of a quantum algorithm with variable stopping times. We want to capture the intuitive idea that an algorithm $\mathcal{A}$ has $m$ potential stopping times $t_1, t_2, \ldots, t_m$. At time $t_j$, the algorithm can indicate that it wants to stop by setting the $j^{\text{th}}$ qubit of a special clock register to 1, where the algorithm starts with all clock qubits set to 0. To enforce that the algorithm actually halts when it indicates that it has, we require that subsequent operations do not affect branches of the computation that have halted in previous steps. We formalize this in the following definition.

**Definition 20** (Variable-time quantum algorithm)**.** *Let $\mathcal{A}$ be a quantum algorithm on a space $\mathcal{H}$ that starts in the state $|0\rangle_{\mathcal{H}}$, the all zeros state in $\mathcal{H}$. We say $\mathcal{A}$ is a* variable-time quantum algorithm *if the following conditions hold:*

1. *$\mathcal{A}$ can be written as the product of $m$ algorithms, $\mathcal{A} = \mathcal{A}_m\mathcal{A}_{m-1}\cdots\mathcal{A}_1$.*
2. *$\mathcal{H}$ can be written as a product $\mathcal{H} = \mathcal{H}_C \otimes \mathcal{H}_A$, where $\mathcal{H}_C$ is a product of $m$ single qubit registers denoted $\mathcal{H}_{C_1}, \mathcal{H}_{C_2}, \ldots, \mathcal{H}_{C_m}$.*
3. *Each $\mathcal{A}_j$ is a controlled unitary that acts on the registers $\mathcal{H}_{C_j} \otimes \mathcal{H}_A$ controlled on the first $j-1$ qubits of $\mathcal{H}_C$ being set to 0.*

In this definition the $m$ potential stopping times of $\mathcal{A}$ correspond to when the segments $\mathcal{A}_j$ end. The last condition enforces that the algorithm in step $j$ does not perform any operation in branches of the computation that have halted in previous steps.

Now consider a variable-time quantum algorithm that prepares a state $|\psi_{\text{succ}}\rangle$ probabilistically. More precisely, the algorithm has a single-qubit flag register that is measured at the end of the algorithm: if we observe the outcome 1 on measuring this register, we have successfully prepared the

state $|\psi_{\text{succ}}\rangle$. Let $p_{\text{succ}}$ denote the probability of obtaining the desired outcome. If the algorithm's complexity is $t_m$, then simply repeating the algorithm $O(1/p_{\text{succ}})$ times will yield an algorithm that creates the state $|\psi_{\text{succ}}\rangle$ with high probability and that has complexity $O(t_m/p_{\text{succ}})$. Using standard amplitude amplification, we can do the same with complexity only $O(t_m/\sqrt{p_{\text{succ}}})$. Variable-time amplitude amplification now allows us to achieve the same with even lower cost if the average stopping time of the algorithm is smaller than its maximum running time, as shown by Ambainis [Amb12]. We specialize the result to our application below.

**Theorem 21** (Variable-time amplitude amplification)**.** *Let $\mathcal{A} = \mathcal{A}_m\mathcal{A}_{m-1}\cdots\mathcal{A}_1$ be a variable-time quantum algorithm on the space $\mathcal{H} = \mathcal{H}_C \otimes \mathcal{H}_F \otimes \mathcal{H}_W$. If $|0\rangle_{\mathcal{H}}$ denotes the all zeros state in $\mathcal{H}$ and $t_j$ denotes the query complexity of the algorithm $\mathcal{A}_j\mathcal{A}_{j-1}\cdots\mathcal{A}_1$, we define*

$$p_j = \|\Pi_{C_j}\mathcal{A}_j\cdots\mathcal{A}_1|0\rangle_{\mathcal{H}}\|^2 \qquad \text{and} \qquad t_{\text{avg}} = \sqrt{\sum_{j=1}^{m} p_j t_j^2} \tag{87}$$

*to be the probability of halting at step $j$ and the average query complexity of the algorithm, respectively, where $\Pi_{C_j}$ denotes the projector onto $|1\rangle$ in $\mathcal{H}_{C_j}$. Additionally, let the success probability of the algorithm and the corresponding output state be denoted*

$$p_{\text{succ}} = \|\Pi_F\mathcal{A}_m\cdots\mathcal{A}_1|0\rangle_{\mathcal{H}}\|^2 \quad \text{and} \quad |\psi_{\text{succ}}\rangle = \frac{\Pi_F\mathcal{A}_m\cdots\mathcal{A}_1|0\rangle_{\mathcal{H}}}{\|\Pi_F\mathcal{A}_m\cdots\mathcal{A}_1|0\rangle_{\mathcal{H}}\|}, \tag{88}$$

*where $\Pi_F$ projects onto $|1\rangle$ in $\mathcal{H}_F$. Then there exists a quantum algorithm with query complexity*

$$O\left(\left(t_m + \frac{t_{\text{avg}}}{\sqrt{p_{\text{succ}}}}\right)\text{poly}(\log t_m)\right) \tag{89}$$

*that produces the state $|\psi_{\text{succ}}\rangle$ with high probability, and outputs a bit indicating whether it was successful. The resulting algorithm is gate-efficient if the algorithm $\mathcal{A}$ is gate-efficient.*

**Gapped phase estimation** Our algorithm also requires the following simple variant of phase estimation, where the goal is to determine whether an eigenphase $\theta \in [-1, 1]$ of a unitary $U$ satisfies $0 \le |\theta| \le \varphi$ or $2\varphi \le |\theta| \le 1$.

**Lemma 22** (Gapped phase estimation)**.** *Let $U$ be a unitary operation with eigenvectors $|\theta\rangle$ satisfying $U|\theta\rangle = e^{i\theta}|\theta\rangle$, and assume that $\theta \in [-1, 1]$. Let $\varphi \in (0, 1/4]$ and let $\epsilon > 0$. Then there is a unitary procedure $\text{GPE}(\varphi, \epsilon)$ making $O(\frac{1}{\varphi}\log\frac{1}{\epsilon})$ queries to $U$ that, on input $|0\rangle_C|0\rangle_P|\theta\rangle$, prepares a state $(\beta_0|0\rangle_C|\gamma_0\rangle_P + \beta_1|1\rangle_C|\gamma_1\rangle_P)|\theta\rangle$, where $|\beta_0|^2 + |\beta_1|^2 = 1$, such that*

- *if $0 \le |\theta| \le \varphi$ then $|\beta_1| \le \epsilon$ and*
- *if $2\varphi \le |\theta| \le 1$ then $|\beta_0| \le \epsilon$.*

*Here $C$ and $P$ are registers of 1 and $\ell$ qubits, respectively, where $\ell = O(\log(1/\varphi)\log(1/\epsilon))$.*

*Proof.* Standard phase estimation gives an estimate of $\theta$ with precision $\varphi/2$ (which suffices to distinguish between the two cases), succeeding with constant probability [CEMM98], using $O(1/\varphi)$ queries to $U$ and $O(\log(1/\varphi))$ ancillary qubits for the estimate of the phase. To boost the success probability, we simply repeat the procedure $O(\log(1/\epsilon))$ times. We construct $\text{GPE}(\varphi, \epsilon)$ by taking a majority vote and encoding the result in the register $C$. By a standard Chernoff bound, this suffices to ensure error probability at most $\epsilon$. □

**Approximating $A^{-1}$** Our algorithm $\mathcal{A}$ also requires a procedure to implement various approximations of $A^{-1}$ that are accurate for different ranges of eigenvectors. If we are promised that the input state $|\psi\rangle$ is a superposition of eigenstates of $A$ with eigenvalues in the range $[1,-\lambda]\cup[\lambda,1]$ for some given $\lambda\in(0,1]$, we can use the results of Section 3 or Section 4 to implement a version of $A^{-1}$ accurate for this range of eigenvalues more efficiently than in the general case.

**Lemma 23.** *For any $\delta>0$ and $\lambda>0$, there exists a unitary $W(\lambda,\delta)$ with query complexity $O(\frac{d}{\lambda}\log^2(\frac{d\kappa}{\delta}))$ satisfying*

$$W(\lambda,\delta)\,|0\rangle_F\,|0\rangle_Q\,|\psi\rangle_I = \frac{1}{\alpha_{\max}}\,|1\rangle_F\,|0\rangle_Q\,h(A)\,|\psi\rangle_I + |0\rangle_F\,|\Psi^\perp\rangle_{QI} \tag{90}$$

*where $\alpha_{\max}=\widetilde{O}(\kappa)$ is a constant independent of $\lambda$, $|\Psi^\perp\rangle_{QI}$ is an unnormalized quantum state on registers $Q$ and $I$ orthogonal to $|0\rangle_Q$, and $\|h(A)|\psi\rangle - A^{-1}|\psi\rangle\|\le\delta$ for any $|\psi\rangle$ that is a superposition of eigenstates of $A$ with eigenvalues in the range $[-1,-\lambda]\cup[\lambda,1]$. Here $F$ and $I$ are registers of 1 and $\log_2 N$ qubits, respectively.*

*Proof.* Using either the Fourier or Chebyshev approach, by choosing an appropriately accurate series expansion of $A^{-1}$ we can perform the map

$$|0\rangle_F\,|0\rangle_Q\,|\psi\rangle_I \mapsto \frac{1}{\alpha}\,|1\rangle_F\,|0\rangle_Q\,h(A)\,|\psi\rangle_I + |0\rangle_F\,|\Psi^\perp\rangle_{QI}\,. \tag{91}$$

Using the Fourier approach, the query complexity of the map in (91) is $O((d/\lambda)\log^2(1/\lambda\delta))$. Using the Chebyshev approach, the query complexity is $O((d/\lambda)\log(d/\lambda\delta))$. In either case, the complexity of the operation in (91) is $O(\frac{d}{\lambda}\log^2(\frac{d\kappa}{\delta}))$.

However, the value of $\alpha$ in (91) will depend on $\lambda$. The constant $\alpha$ can be determined from the $L_1$ norm in the approximation of $A^{-1}$ by a sum of unitaries. Under the promise that the input state is a superposition of eigenvectors with eigenvalues in $[-1,-\lambda]\cup[\lambda,1]$, the value of $\alpha$ depends on $\lambda$ and $\delta$, so we denote it $\alpha=\alpha(\lambda,\delta)$.

Since we require transformations for which the factor $1/\alpha$ appearing in (91) is the same independent of $\lambda$, we choose the largest value of $\alpha$ for our range of eigenvalues, which is $\alpha_{\max}:=\alpha(1/\kappa,\delta)$. This value is $\alpha_{\max}=O(\kappa\sqrt{\log(\kappa/\delta)})$ or $\alpha_{\max}=O(\kappa\log(d\kappa/\delta))$ for the Fourier or Chebyshev approaches, respectively, which is $\widetilde{O}(\kappa)$ in either case.

Since $\alpha(\lambda,\delta)\le\alpha_{\max}$, and $\alpha(\lambda,\delta)$ can be computed efficiently, if we implement (91) with a smaller value of $\alpha$ than needed, it is always possible to make the implementation worse and increase the value of $\alpha$. In particular, the algorithm for preparing (91) with $\alpha$ replaced by the fixed value $\alpha_{\max}$ is the algorithm described in Section 3 or Section 4 (which implements (91)) followed by an additional step that transforms

$$|1\rangle_F \mapsto \frac{\alpha(\lambda,\delta)}{\alpha_{\max}}\,|1\rangle_F + \sqrt{1-\frac{\alpha(\lambda,\delta)^2}{\alpha_{\max}^2}}\,|0\rangle_F\,, \tag{92}$$

conditional on $|0\rangle_Q$. This gives us a map $W(\lambda,\delta)$ that performs the transformation in (91) with the same value $\alpha_{\max}$ in the denominator independent of $\lambda$, as described in the lemma. □

## 5.2 Algorithm

We now describe a variable-time quantum algorithm $\mathcal{A}$ built as a sequence of steps $\mathcal{A}_1,\dots,\mathcal{A}_m$ with $m:=\lceil\log_2\kappa\rceil+1$, so the algorithm is $\mathcal{A}=\mathcal{A}_m\cdots\mathcal{A}_1$. The algorithm $\mathcal{A}$ uses the following registers:

- an $m$-qubit clock register $C$, labeled $C_1, \ldots, C_m$, used to determine a region the eigenvalue belongs to (i.e., to store the result of GPE);
- a single-qubit flag register $F$ to indicate whether the approximation of $A^{-1}$ was successfully implemented;
- a $(\log_2 N)$-qubit input register $I$, initialized to $|b\rangle$, that finally contains the output state;
- a register $P$, divided into registers $P_1, P_2, \ldots, P_m$, to be used as ancilla for GPE; and
- a register $Q$ to be used as ancilla in the implementation of $A^{-1}$.

The corresponding Hilbert spaces are denoted $\mathcal{H}_C$, $\mathcal{H}_F$, $\mathcal{H}_I$, $\mathcal{H}_P$, and $\mathcal{H}_Q$, respectively. All registers are initialized in $|0\rangle$ except for register $I$, which is initialized in $|b\rangle$. When we write $|0\rangle_X$ we mean that all qubits of register $X$ are in $|0\rangle$.

**Algorithm $\mathcal{A}_j$** We now describe the algorithm $\mathcal{A}_j$, which forms a part of the variable-time algorithm $\mathcal{A}$. In the algorithm below, each call to GPE (Lemma 22) uses the unitary operation $U := e^{iA}$. Note that this unitary satisfies the assumption of Lemma 22 as the norm of $A$ is at most 1. For all $j \in [m]$, let $\varphi_j := 2^{-j}$, and let $\delta = \epsilon/(m\alpha_{\max})$.

Finally, we define $\mathcal{A}_j$ as the product of the following two unitary operations:

1. Conditional on first $j-1$ qubits of $\mathcal{H}_C$ being $|0\rangle$, apply GPE$(\varphi_j, \delta)$ on the input state in $I$ using $C_j$ as the output qubit and additional fresh qubits from $P$ as ancilla (denoted $P_j$).
2. Conditional on $C_j$ (the outcome of the previous step) being $|1\rangle_{C_j}$, apply $W(\varphi_j, m\delta)$ on the input state in $I$ using $F$ as the flag register and register $Q$ as ancilla.

**Final algorithm** Before describing the final algorithm, we define another sequence of algorithms $\mathcal{A}' = \mathcal{A}'_m \cdots \mathcal{A}'_1$, similar to $\mathcal{A}$. The only difference between $\mathcal{A}_j$ and $\mathcal{A}'_j$ is that instead of applying the operator $W$ in step 2 of $\mathcal{A}_j$, algorithm $\mathcal{A}'_j$ applies the following trivial operation $W'$:

$$W' |0\rangle_F |0\rangle_Q |\psi\rangle_I = |1\rangle_F |0\rangle_Q |\psi\rangle_I . \tag{93}$$

The main algorithm of this section first applies VTAA to $\mathcal{A}$ to produce a normalized version of the state output by $\mathcal{A}$ that has register $F$ in $|1\rangle_F$. The algorithm then applies the unitary $(\mathcal{A}')^\dagger$ to this state. At this point the state in register $I$ will be $\frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|}$ up to error $O(\epsilon)$, as we prove in the next section.

### 5.3 Correctness

To prove that the algorithm works correctly, it is useful to first analyze its behavior on an eigenstate $|\lambda\rangle$ of $A$, which satisfies $A|\lambda\rangle = \lambda|\lambda\rangle$ for some $\lambda \in [-1, 1]$. Then the action of $\mathcal{A}$ on a general state $|b\rangle$ follows from linearity.

Let $j \in [m]$ be such that $\varphi_j < |\lambda| \leq 2\varphi_j$ (recall that $\varphi_j$ was defined to be $1/2^j$). Such a $j$ must exist because the largest value of $|\lambda|$ is 1, which is at most $2\varphi_1$, and the smallest value of $|\lambda|$ is $1/\kappa$, which satisfies $\varphi_m = 1/2^{\lceil \log(\kappa) \rceil + 1} < 1/\kappa$.

**State after $\mathcal{A}_1$ to $\mathcal{A}_{j-1}$** We first observe that the algorithms $\mathcal{A}_1, \cdots, \mathcal{A}_{j-1}$ essentially do nothing to the state except modify the ancilla register $P$. To see this, observe that we start with the state $|0\rangle_C|0\rangle_F|\lambda\rangle_I|0\rangle_P|0\rangle_Q$. Upon applying step 1 of $\mathcal{A}_1$, because $|\lambda| \leq \phi_1$, the clock register $C_1$ remains $|0\rangle$ with very high probability, and only the register $P_1$ has been modified to hold the ancillary state of GPE. Thus, we have the state

$$|0\rangle_C|0\rangle_F|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}|0\rangle_{P_2 \cdots P_m}|0\rangle_Q \tag{94}$$

up to error $O(\delta)$, where $|\gamma_0^1\rangle$ is the ancillary state produced by GPE. Step 2 of $\mathcal{A}_1$ does nothing because register $C_1$ is set to $|0\rangle$. Similarly, after $j-1$ steps we are left with the state

$$|0\rangle_C|0\rangle_F|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^{j-1}\rangle_{P_{j-1}}|0\rangle_{P_j\cdots P_m}|0\rangle_Q \tag{95}$$

up to error $O(j\delta) = O(m\delta)$, where $|\gamma_0^i\rangle$ is the ancillary state produced by the $i^{\text{th}}$ call to GPE.

**State after $\mathcal{A}_j$** Now when we apply step 1 of $\mathcal{A}_j$, because $|\lambda|$ lies in between $\varphi_j$ and $2\varphi_j$, GPE will return a superposition of $|0\rangle_{C_j}$ and $|1\rangle_{C_j}$. We will then have the state

$$\beta_0|0\rangle_C|0\rangle_F|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^{j-1}\rangle_{P_{j-1}}|\gamma_0^j\rangle_{P_j}|0\rangle_{P_{j+1}\cdots P_m}|0\rangle_Q \tag{96}$$

$$+\beta_1|\mathfrak{U}_j\rangle_C|0\rangle_F|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^{j-1}\rangle_{P_{j-1}}|\gamma_1^j\rangle_{P_j}|0\rangle_{P_{j+1}\cdots P_m}|0\rangle_Q, \tag{97}$$

where $\mathfrak{U}_j := 0^{j-1}10^{m-j}$ denotes the integer $j$ represented in unary. Now when step 2 of $\mathcal{A}_j$ is applied, the part of the state in (96) remains unchanged, since the register $C_j$ is set to $|0\rangle$. However, the part in (97) will have operator $W(\varphi_j, m\delta)$ applied to it and the state in (97) will become

$$\beta_1|\mathfrak{U}_j\rangle_C|1\rangle_F\Big(\frac{h(A)}{\alpha_{\max}}|\lambda\rangle_I\Big)|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^{j-1}\rangle_{P_{j-1}}|\gamma_1^j\rangle_{P_j}|0\rangle_{P_{j+1}\cdots P_m}|0\rangle_Q + \beta_1|\mathfrak{U}_j\rangle_C|0\rangle_F|g_j\rangle_{IPQ}, \tag{98}$$

where $|g_j\rangle_{IPQ}$ is some state on register $I$, $P$, $Q$. Since $\lambda$ falls in the required range of Lemma 23, i.e., it is between $[-1,-\varphi_j]\cup[\varphi_j,1]$, we can replace $h(A)$ in the equation above with $A^{-1}$, incurring error at most $O(m\delta)$.

**State after $\mathcal{A}_{j+1}$** We can now apply the algorithm $\mathcal{A}_{j+1}$ to our state. This will only affect the part of the state in (96), because the computation in (98) has ended and has a nonzero value in register $C$. On applying step 1 of $\mathcal{A}_{j+1}$, the state (96) is mapped to

$$\beta_0|\mathfrak{U}_{j+1}\rangle_C|0\rangle_F|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^j\rangle_{P_j}|\gamma_1^{j+1}\rangle_{P_{j+1}}|0\rangle_{P_{j+2}\cdots P_m}|0\rangle_Q \tag{99}$$

up to error $O(m\delta)$, since $|\lambda| > 2\varphi_{j+1}$, and hence GPE outputs $|1\rangle$ with very high probability. We can now apply step 2 of $\mathcal{A}_{j+1}$ to this state to obtain

$$\beta_0|\mathfrak{U}_{j+1}\rangle_C|1\rangle_F\Big(\frac{h(A)}{\alpha_{\max}}|\lambda\rangle_I\Big)|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^j\rangle_{P_j}|\gamma_1^{j+1}\rangle_{P_{j+1}}|0\rangle_{P_{j+2}\cdots P_m}|0\rangle_Q + \beta_0|\mathfrak{U}_{j+1}\rangle_C|0\rangle_F|g_{j+1}\rangle_{IPQ} \tag{100}$$

up to error $O(m\delta)$, where $|g_{j+1}\rangle_{IPQ}$ is some state on register $I$, $P$, $Q$. Since $\lambda$ falls in the required range of Lemma 23, i.e., it is between $[-1,-\varphi_{j+1}]\cup[\varphi_{j+1},1]$, we can replace $h(A)$ in the equation above with $A^{-1}$, incurring error at most $O(m\delta)$.

**State after $\mathcal{A}$** Now since the state has no overlap with $|0\rangle_C$, the remaining operations of $\mathcal{A}$ do nothing, and at the end of the algorithm $\mathcal{A}$, we have the state

$$\begin{aligned}
&\frac{\beta_1}{\alpha_{\max}}|\mathfrak{U}_j\rangle_C|1\rangle_F A^{-1}|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^{j-1}\rangle_{P_{j-1}}|\gamma_1^j\rangle_{P_j}|0\rangle_{P_{j+1}\cdots P_m}|0\rangle_Q\\
&+\frac{\beta_0}{\alpha_{\max}}|\mathfrak{U}_{j+1}\rangle_C|1\rangle_F A^{-1}|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^j\rangle_{P_j}|\gamma_1^{j+1}\rangle_{P_{j+1}}|0\rangle_{P_{j+2}\cdots P_m}|0\rangle_Q\\
&+\beta_1|\mathfrak{U}_j\rangle_C|0\rangle_F|g_j\rangle_{IPQ} + \beta_0|\mathfrak{U}_{j+1}\rangle_C|0\rangle_F|g_{j+1}\rangle_{IPQ}
\end{aligned} \tag{101}$$

up to error $O(m\delta)$. On projecting this state to the $|1\rangle_F$ subspace, we obtain

$$\Pi_F \mathcal{A}|0\rangle_{CF}|\lambda\rangle_I|0\rangle_{PQ} = \frac{\beta_1}{\alpha_{\max}}|\mathfrak{U}_j\rangle_C|1\rangle_F A^{-1}|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^{j-1}\rangle_{P_{j-1}}|\gamma_1^j\rangle_{P_j}|0\rangle_{P_{j+1}\cdots P_m}|0\rangle_Q + \frac{\beta_0}{\alpha_{\max}}|\mathfrak{U}_{j+1}\rangle_C|1\rangle_F A^{-1}|\lambda\rangle_I|\gamma_0^1\rangle_{P_1}\cdots|\gamma_0^{j-1}\rangle_{P_{j-1}}|\gamma_0^j\rangle_{P_j}|\gamma_1^{j+1}\rangle_{P_{j+1}}|0\rangle_{P_{j+2}\cdots P_m}|0\rangle_Q, \quad (102)$$

where $\Pi_F$ denotes the projector onto $|1\rangle_F$. Now let $|\Psi_\lambda\rangle$ denote a normalized quantum state on registers $C$, $F$, $P$, and $Q$, such that this equation can be rewritten as

$$\Pi_F \mathcal{A}|0\rangle_{CF}|\lambda\rangle_I|0\rangle_{PQ} = \frac{A^{-1}}{\alpha_{\max}}|\lambda\rangle_I|\Psi_\lambda\rangle_{CFPQ}. \quad (103)$$

**Algorithm $\mathcal{A}'$** Now observe that if we had run algorithm $\mathcal{A}'$ instead of algorithm $\mathcal{A}$, the ancillary state $|\Psi_\lambda\rangle$ created would be identical since this state depends only on GPE and not on the implementation of $A^{-1}$. In $\mathcal{A}'$, since the operation $W$ is replaced by the operation in (93) that simply flips the bit in register $F$, we have

$$\mathcal{A}'|0\rangle_{CF}|\lambda\rangle_I|0\rangle_{PQ} = |\lambda\rangle_I|\Psi_\lambda\rangle_{CFPQ}. \quad (104)$$

Note that there is no need to project onto $|1\rangle_F$, since the final state here has no overlap on $|0\rangle_F$, and hence there is no need to apply VTAA either. Thus the inverse of algorithm $\mathcal{A}'$ can be used to erase the state $|\Psi_\lambda\rangle$ given state $|\lambda\rangle$ in another register, even in superposition.

**Final algorithm** We can now analyze the general case where the input state is $|b\rangle = \sum_k c_k |\lambda_k\rangle$, where $|\lambda_k\rangle$ are eigenvectors of $A$ of eigenvalue $\lambda_k \in [-1, 1]$ and $\sum_k |c_k|^2 = 1$.

Using linearity and (103), we obtain

$$\Pi_F \mathcal{A}|0\rangle_{CF}|b\rangle_I|0\rangle_{PQ} = \frac{A^{-1}}{\alpha_{\max}}\sum_k c_k|\lambda_k\rangle_I|\Psi_{\lambda_k}\rangle_{CFPQ}. \quad (105)$$

We can now apply VTAA (Theorem 21) to $\mathcal{A}$ to produce a normalized version of this state with constant probability. The approximation error is $O(\alpha_{\max} m\delta) = O(\epsilon)$ since, in the worst case, the norm of this state is $O(1/\alpha_{\max})$. Now applying $(\mathcal{A}')^\dagger$ allows us to erase the ancillary states $|\Psi_{\lambda_k}\rangle_{CFPQ}$ to obtain

$$\frac{A^{-1}|b\rangle_I|0\rangle_{CFPQ}}{\|A^{-1}|b\rangle_I|0\rangle_{CFPQ}\|} \quad (106)$$

up to error $O(\epsilon)$, as desired.

### 5.4 Complexity

The complexity of our algorithm is dominated by the cost of applying VTAA to algorithm $\mathcal{A}$. The next step, which is to apply $(\mathcal{A}')^\dagger$, can only cost as much as $\mathcal{A}$, and hence can be ignored.

To obtain the query complexity of applying VTAA to $\mathcal{A}$, we need to compute the quantities in Theorem 21, which are $t_m$, $p_{\text{succ}}$, and $t_{\text{avg}}$.

**Computing $t_j$** Let us begin with $t_j$, the query complexity of applying $\mathcal{A}_j \cdots \mathcal{A}_1$. The cost of any $\mathcal{A}_j$ is the cost of applying $\mathrm{GPE}(\varphi_j, \delta)$ using Lemma 22 and $W(\varphi_j, m\delta)$ using Lemma 23. The number of uses of $U = e^{iA}$ in $\mathrm{GPE}(\varphi_j, \delta)$ follows from Lemma 22 and is $O(2^j \log(1/\delta))$. Then, each $U$ must be implemented using a Hamiltonian simulation algorithm within precision $O(\delta/(2^j \log(1/\delta)))$. By the results of [BCK15], the query complexity of $\mathrm{GPE}(\varphi_j, \delta)$ is therefore

$$O(d2^j \operatorname{poly}(\log(d2^j/\delta))) = \widetilde{O}(d2^j). \tag{107}$$

The query complexity of $W(\varphi_j, m\delta)$ is $O(\frac{d}{\varphi_j} \log^2(\frac{d\kappa}{m\delta})) = \widetilde{O}(d2^j)$. Since the cost of $\mathcal{A}_j$ is $\widetilde{O}(d2^j)$, $t_j = \widetilde{O}(d2^j)$ and $t_m = \widetilde{O}(d\kappa)$.

**Computing $p_{\mathrm{succ}}$** The probability $p_{\mathrm{succ}}$ is the probability of measuring the register $F$ in $|1\rangle_F$ in the state output by algorithm $\mathcal{A}$. This is simply the squared norm of the state in (105), which is

$$p_{\mathrm{succ}} = \left\| \frac{A^{-1}}{\alpha_{\max}} \sum_k c_k |\lambda_k\rangle_I |\Psi_\lambda\rangle_{CFPQ} \right\|^2 = \frac{1}{\alpha_{\max}^2} \left\| \sum_k \frac{c_k}{\lambda_k} |\lambda_k\rangle_I |\Psi_\lambda\rangle_{CFPQ} \right\|^2 = \frac{1}{\alpha_{\max}^2} \sum_k \frac{|c_k|^2}{\lambda_k^2}. \tag{108}$$

To be precise, this is only an approximation of $p_{\mathrm{succ}}$ up to error $O(m\delta) = O(\epsilon/\alpha_{\max})$. However, since this error is much smaller than the calculated value of $p_{\mathrm{succ}}$, this error only affects the value by a constant factor.

**Computing $t_{\mathrm{avg}}$** Let $p_j$ denote the probability that algorithm $\mathcal{A}$ stops exactly at the $j^{\mathrm{th}}$ step, which is $\|\Pi_{C_j} \mathcal{A}_j \cdots \mathcal{A}_1 |b\rangle_I |0\rangle_{CFPQ}\|^2$, where $\Pi_{C_j}$ is the projector onto $|1\rangle_{C_j}$. We can now compute

$$t_{\mathrm{avg}}^2 = \sum_j p_j t_j^2 = \sum_j \|\Pi_{C_j} \mathcal{A}_j \cdots \mathcal{A}_1 |b\rangle_I |0\rangle_{CFPQ}\|^2 t_j^2 \tag{109}$$

$$= \sum_j \|\Pi_{C_j} \mathcal{A}_j \cdots \mathcal{A}_1 \sum_k c_k |\lambda_k\rangle_I |0\rangle_{CFPQ}\|^2 t_j^2 \tag{110}$$

$$= \sum_j \sum_k |c_k|^2 \|\Pi_{C_j} \mathcal{A}_j \cdots \mathcal{A}_1 |\lambda_k\rangle_I |0\rangle_{CFPQ}\|^2 t_j^2 \tag{111}$$

$$= \sum_k |c_k|^2 \Big( \sum_j \|\Pi_{C_j} \mathcal{A}_j \cdots \mathcal{A}_1 |\lambda_k\rangle_I |0\rangle_{CFPQ}\|^2 t_j^2 \Big), \tag{112}$$

where the expression in parentheses is the average stopping time of the algorithm for the special case where $|b\rangle = |\lambda_k\rangle$. For the state $|\lambda_k\rangle$, we know the algorithm stops after step $j$ or $j+1$, where $j$ satisfies $\varphi_j < |\lambda_k| \leq 2\varphi_j$. Hence the average stopping time for $|\lambda_k\rangle$ is at most $t_{j+1} = \widetilde{O}(d2^j) = \widetilde{O}(d/\lambda_k)$. Substituting this into (112), we have

$$t_{\mathrm{avg}}^2 = \sum_k |c_k|^2 \widetilde{O}(d^2/\lambda_k^2) = \widetilde{O}\left( d^2 \sum_k \frac{|c_k|^2}{\lambda_k^2} \right). \tag{113}$$

**Total complexity** We are now ready to compute the total cost of VTAA and prove Theorem 5. Theorem 21 states that the query complexity of VTAA applied to $\mathcal{A}$ is

$$O\left( \left( t_m + \frac{t_{\mathrm{avg}}}{\sqrt{p_{\mathrm{succ}}}} \right) \operatorname{poly}(\log t_m) \right). \tag{114}$$

Using the values we computed, this is

$$\widetilde{O}(d\kappa + d\alpha_{\max}) = \widetilde{O}(d\kappa), \tag{115}$$

where we used the fact that $\alpha_{\max} = \widetilde{O}(\kappa)$, as stated in Lemma 23.

Finally, since $\mathcal{A}$ and $\mathcal{A}'$ are gate-efficient and VTAA preserves this property, the overall algorithm is gate-efficient.

## Acknowledgements

RK thanks Vamsi Pritham Pingali for several helpful discussions during the course of this work. The authors acknowledge support from AFOSR grant number FA9550-12-1-0057, ARO grant numbers W911NF-12-1-0482 and W911NF-12-1-0486, CIFAR, IARPA grant number D15PC00242, NRO, and NSF grant number 1526380. This preprint is MIT-CTP #4687 and LA-UR-15-27205.

## References

[Amb12] Andris Ambainis. Variable time amplitude amplification and quantum algorithms for linear algebra problems. In *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science*, pages 636–647, 2012. arXiv:1010.4458. [pp. 4, 5, 20, 21]

[BC12] Dominic W. Berry and Andrew M. Childs. Black-box Hamiltonian simulation and unitary implementation. *Quantum Information and Computation*, 12(1–2):29–62, 2012. arXiv:0910.4157. [pp. 2, 5, 14, 15, 16]

[BCC+14] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Exponential improvement in precision for simulating sparse Hamiltonians. In *Proceedings of the 46th ACM Symposium on Theory of Computing*, pages 283–292, 2014. [pp. 1, 4, 5]

[BCC+15] Dominic W. Berry, Andrew M. Childs, Richard Cleve, Robin Kothari, and Rolando D. Somma. Simulating Hamiltonian dynamics with a truncated Taylor series. *Physical Review Letters*, 114:090502, 2015. [pp. 4, 5]

[BCK15] Dominic W. Berry, Andrew M. Childs, and Robin Kothari. Hamiltonian simulation with nearly optimal dependence on all parameters. In *Proceedings of the 56th Symposium on Foundations of Computer Science*, 2015. arXiv:1501.01715. [pp. 2, 3, 4, 5, 9, 11, 13, 14, 15, 16, 19, 26]

[BHMT02] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation. In *Quantum computation and information*, volume 305 of *Contemporary Mathematics*, pages 53–74. AMS, 2002. [pp. 6, 7]

[CEMM98] Richard Cleve, Artur Ekert, Chiara Macchiavello, and Michele Mosca. Quantum algorithms revisited. *Proceedings of the Royal Society of London, Series A*, 454:339–354, 1998. [p. 21]

[Chi10] Andrew M. Childs. On the relationship between continuous- and discrete-time quantum walk. *Communications in Mathematical Physics*, 294(2):581–603, 2010. arXiv:0810.0312. [pp. 5, 14]

[CW12] Andrew M. Childs and Nathan Wiebe. Hamiltonian simulation using linear combinations of unitary operations. *Quantum Information and Computation*, 12:901–924, 2012. [p. 4]

[Gro96] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th ACM Symposium on Theory of Computing*, pages 212–219, 1996. [p. 2]

[Har15] Aram W. Harrow. Review of quantum algorithms for systems of linear equations, 2015. arXiv:1501.00008. [p. 2]

[HHL09] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15):150502, 2009. arXiv:0811.3171. [pp. 1, 2, 3, 4]

[Kot14] Robin Kothari. *Efficient algorithms in quantum query complexity*. PhD thesis, University of Waterloo, 2014. [pp. 4, 6]

[SBM06] Vivek V. Shende, Stephen S. Bullock, and Igor L. Markov. Synthesis of quantum-logic circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 25(6):1000–1010, 2006. arXiv:quant-ph/0406176. [pp. 13, 19]

[SOG+02] Rolando D. Somma, Gerardo Ortiz, James E. Gubernatis, Emanuel Knill, and Raymond Laflamme. Simulating physical phenomena by quantum networks. *Physical Review A*, 65:042323, 2002. [p. 4]